

Benchmarking benchmarking,
and optimizing optimization

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Bit operations per bit of plaintext
(assuming precomputed subkeys),
as listed in recent Skinny paper:

key	ops/bit	cipher
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Benchmarking benchmarking,
and optimizing optimization

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

marking benchmarking,
mizing optimization

. Bernstein

ty of Illinois at Chicago &
he Universiteit Eindhoven

1

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

2

Operatio
poor mo
worse m

Pick a c
How fas

First ste
Write sim

e.g. Ber

Janssen-
Smetsers

includes
impleme

enchmarking,
timization

is at Chicago &
iteit Eindhoven

1

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

2

Operation counts
poor model of hardware
worse model of software

Pick a cipher: e.g.
How fast is Salsa20

First step in analysis
Write simple software

e.g. Bernstein–van
Janssen–Lange–Sch
Smetsers “TweetN
includes essentially
implementation of

g,

ago &
hoven

1

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

2

Operation counts are a
poor model of hardware cost
worse model of software cost

Pick a cipher: e.g., Salsa20.
How fast is Salsa20 software

First step in analysis:
Write simple software.

e.g. Bernstein–van Gastel–
Janssen–Lange–Schwabe–
Smetsers “TweetNaCl”
includes essentially the follow
implementation of Salsa20:

Bit operations per bit of plaintext
(assuming precomputed subkeys),
not entirely listed in Skinny paper:

key	ops/bit	cipher
256	54	Salsa20/8
256	78	Salsa20/12
128	88	Simon: 60 ops broken
128	100	NOEKEON
128	117	Skinny
256	126	Salsa20
256	144	Simon: 106 ops broken
128	147.2	PRESENT
256	156	Skinny
128	162.75	Piccolo
128	202.5	AES
256	283.5	AES

Operation counts are a
poor model of hardware cost,
worse model of software cost.

Pick a cipher: e.g., Salsa20.

How fast is Salsa20 software?

First step in analysis:

Write simple software.

e.g. Bernstein–van Gastel–
Janssen–Lange–Schwabe–
Smetsers “TweetNaCl”

includes essentially the following
implementation of Salsa20:

operations per bit of plaintext
(using precomputed subkeys),
are listed in Skinny paper:

ops/bit	cipher
4	Salsa20/8
3	Salsa20/12
3	Simon: 60 ops broken
0	NOEKEON
7	Skinny
5	Salsa20
4	Simon: 106 ops broken
7.2	PRESENT
5	Skinny
2.75	Piccolo
2.5	AES
3.5	AES

2

Operation counts are a
poor model of hardware cost,
worse model of software cost.

Pick a cipher: e.g., Salsa20.

How fast is Salsa20 software?

First step in analysis:

Write simple software.

e.g. Bernstein–van Gastel–
Janssen–Lange–Schwabe–
Smetsers “TweetNaCl”

includes essentially the following
implementation of Salsa20:

3

```
int crypto
const u8 *
{
    u32 w[16]
    int i, j

    FOR(i, 4)
        x[5+i]
        x[1+i]
        x[6+i]
        x[11+i]
    }
    FOR(i, 16)
```


2

bit of plaintext
 (computed subkeys),
 in Skinny paper:

er

a20/8

a20/12

on: 60 ops broken

EKEON

ny

a20

on: 106 ops broken

ESSENT

ny

olo

Operation counts are a
 poor model of hardware cost,
 worse model of software cost.

Pick a cipher: e.g., Salsa20.

How fast is Salsa20 software?

First step in analysis:

Write simple software.

e.g. Bernstein–van Gastel–

Janssen–Lange–Schwabe–

Smetsers “TweetNaCl”

includes essentially the following

implementation of Salsa20:

3

```
int crypto_core_salsa
const u8 *in,const u8
{
    u32 w[16],x[16],y[16];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4);
        x[1+i] = ld32(k+4);
        x[6+i] = ld32(in+4);
        x[11+i] = ld32(k+4);
    }

    FOR(i,16) y[i] = x[i]
```


2

intext
(keys),
paper:

Operation counts are a poor model of hardware cost, worse model of software cost.

Pick a cipher: e.g., Salsa20.

How fast is Salsa20 software?

First step in analysis:

Write simple software.

e.g. Bernstein–van Gastel–

Janssen–Lange–Schwabe–

Smetsers “TweetNaCl”

includes essentially the following implementation of Salsa20:

3

```
int crypto_core_salsa20(u8 *out,
const u8 *in,const u8 *k,const u
{
    u32 w[16],x[16],y[16],t[4];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4*i);
        x[1+i] = ld32(k+4*i);
        x[6+i] = ld32(in+4*i);
        x[11+i] = ld32(k+16+4*i);
    }

    FOR(i,16) y[i] = x[i];
```

broken

os broken

Operation counts are a poor model of hardware cost, worse model of software cost.

Pick a cipher: e.g., Salsa20.

How fast is Salsa20 software?

First step in analysis:

Write simple software.

e.g. Bernstein–van Gastel–

Janssen–Lange–Schwabe–

Smetsers “TweetNaCl”

includes essentially the following

implementation of Salsa20:

```
int crypto_core_salsa20(u8 *out,
const u8 *in,const u8 *k,const u8 *c)
{
    u32 w[16],x[16],y[16],t[4];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4*i);
        x[1+i] = ld32(k+4*i);
        x[6+i] = ld32(in+4*i);
        x[11+i] = ld32(k+16+4*i);
    }

    FOR(i,16) y[i] = x[i];
```

on counts are a
odel of hardware cost,
odel of software cost.

ipher: e.g., Salsa20.
t is Salsa20 software?

p in analysis:
mple software.

nstein–van Gastel–
–Lange–Schwabe–
s “TweetNaCl”

essentially the following
ntation of Salsa20:

3

```
int crypto_core_salsa20(u8 *out,  
const u8 *in,const u8 *k,const u8 *c)  
{  
    u32 w[16],x[16],y[16],t[4];  
    int i,j,m;  
  
    FOR(i,4) {  
        x[5*i] = ld32(c+4*i);  
        x[1+i] = ld32(k+4*i);  
        x[6+i] = ld32(in+4*i);  
        x[11+i] = ld32(k+16+4*i);  
    }  
  
    FOR(i,16) y[i] = x[i];  
}
```

4

```
FOR(i,20)  
    FOR(j,  
        FOR(  
            t[1]  
            t[2]  
            t[3]  
            t[0]  
            FOR(  
        }  
        FOR(m  
    }  
  
    FOR(i,16)  
    return (  
}
```

3

```

int crypto_core_salsa20(u8 *out,
const u8 *in,const u8 *k,const u8 *c)
{
    u32 w[16],x[16],y[16],t[4];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4*i);
        x[1+i] = ld32(k+4*i);
        x[6+i] = ld32(in+4*i);
        x[11+i] = ld32(k+16+4*i);
    }

    FOR(i,16) y[i] = x[i];

```

4

```

    FOR(i,20) {
        FOR(j,4) {
            FOR(m,4) t[m] =
            t[1] ^= L32(t[0]);
            t[2] ^= L32(t[1]);
            t[3] ^= L32(t[2]);
            t[0] ^= L32(t[3]);
            FOR(m,4) w[4*j+m] =
        }
        FOR(m,16) x[m] =
    }

    FOR(i,16) st32(out+i,y[i]);
    return 0;
}

```

3

```

int crypto_core_salsa20(u8 *out,
const u8 *in,const u8 *k,const u8 *c)
{
    u32 w[16],x[16],y[16],t[4];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4*i);
        x[1+i] = ld32(k+4*i);
        x[6+i] = ld32(in+4*i);
        x[11+i] = ld32(k+16+4*i);
    }

    FOR(i,16) y[i] = x[i];

```

4

```

    FOR(i,20) {
        FOR(j,4) {
            FOR(m,4) t[m] = x[(5*j+4*m)];
            t[1] ^= L32(t[0]+t[3], 7);
            t[2] ^= L32(t[1]+t[0], 9);
            t[3] ^= L32(t[2]+t[1], 13);
            t[0] ^= L32(t[3]+t[2], 18);
            FOR(m,4) w[4*j+(j+m)%4] =
        }
        FOR(m,16) x[m] = w[m];
    }

    FOR(i,16) st32(out + 4 * i,x[i]);
    return 0;
}

```

```

int crypto_core_salsa20(u8 *out,
const u8 *in,const u8 *k,const u8 *c)
{
    u32 w[16],x[16],y[16],t[4];
    int i,j,m;

    FOR(i,4) {
        x[5*i] = ld32(c+4*i);
        x[1+i] = ld32(k+4*i);
        x[6+i] = ld32(in+4*i);
        x[11+i] = ld32(k+16+4*i);
    }

    FOR(i,16) y[i] = x[i];

```

```

    FOR(i,20) {
        FOR(j,4) {
            FOR(m,4) t[m] = x[(5*j+4*m)%16];
            t[1] ^= L32(t[0]+t[3], 7);
            t[2] ^= L32(t[1]+t[0], 9);
            t[3] ^= L32(t[2]+t[1], 13);
            t[0] ^= L32(t[3]+t[2], 18);
            FOR(m,4) w[4*j+(j+m)%4] = t[m];
        }
        FOR(m,16) x[m] = w[m];
    }

    FOR(i,16) st32(out + 4 * i,x[i] + y[i]);
    return 0;
}

```

4

```

o_core_salsa20(u8 *out,
*in,const u8 *k,const u8 *c)
6],x[16],y[16],t[4];
,m;
) {
] = ld32(c+4*i);
] = ld32(k+4*i);
] = ld32(in+4*i);
i] = ld32(k+16+4*i);
6) y[i] = x[i];

```

```

FOR(i,20) {
    FOR(j,4) {
        FOR(m,4) t[m] = x[(5*j+4*m)%16];
        t[1] ^= L32(t[0]+t[3], 7);
        t[2] ^= L32(t[1]+t[0], 9);
        t[3] ^= L32(t[2]+t[1],13);
        t[0] ^= L32(t[3]+t[2],18);
        FOR(m,4) w[4*j+(j+m)%4] = t[m];
    }
    FOR(m,16) x[m] = w[m];
}
FOR(i,16) st32(out + 4 * i,x[i] + y[i]);
return 0;
}

```

5

```

static con
= "expand
int crypt
const u8 *
{
    u8 z[16]
    u32 u,i
    if (!b)
    FOR(i,16)
    FOR(i,8)
    while (1)
        crypt
    FOR(i
    u = 1

```


4

```

a20(u8 *out,
    u8 *k, const u8 *c)
{
    u32 t[4];
    FOR(i, 20) {
        t[i] = L32(x[i] + y[i]);
        FOR(j, 4) {
            t[j] ^= L32(t[(j+1)%4] + t[(j+2)%4], sigma[j]);
        }
    }
    FOR(i, 16) st32(out + 4 * i, x[i] + y[i]);
    return 0;
}

```

```

FOR(i, 20) {
    FOR(j, 4) {
        FOR(m, 4) t[m] = x[(5*j+4*m)%16];
        t[1] ^= L32(t[0]+t[3], 7);
        t[2] ^= L32(t[1]+t[0], 9);
        t[3] ^= L32(t[2]+t[1], 13);
        t[0] ^= L32(t[3]+t[2], 18);
        FOR(m, 4) w[4*j+(j+m)%4] = t[m];
    }
    FOR(m, 16) x[m] = w[m];
}

FOR(i, 16) st32(out + 4 * i, x[i] + y[i]);
return 0;
}

```

5

```

static const u8 sigma[4] = {
    "expand 32-byte k";
};

int crypto_stream_salsa8(u8 *c, u64 b, const u8 *m)
{
    u8 z[16], x[64];
    u32 u, i;
    if (!b) return 0;
    FOR(i, 16) z[i] = 0;
    FOR(i, 8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa8(u, z, x);
        FOR(i, 64) c[i] = m[i] ^ x[i];
        u = 1;
    }
}

```

4

8 *c)

```

FOR(i,20) {
    FOR(j,4) {
        FOR(m,4) t[m] = x[(5*j+4*m)%16];
        t[1] ^= L32(t[0]+t[3], 7);
        t[2] ^= L32(t[1]+t[0], 9);
        t[3] ^= L32(t[2]+t[1], 13);
        t[0] ^= L32(t[3]+t[2], 18);
        FOR(m,4) w[4*j+(j+m)%4] = t[m];
    }
    FOR(m,16) x[m] = w[m];
}

FOR(i,16) st32(out + 4 * i, x[i] + y[i]);
return 0;
}

```

5

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8
const u8 *m,u64 b,const u8 *n,co
{
    u8 z[16],x[64];
    u32 u,i;
    if (!b) return 0;
    FOR(i,16) z[i] = 0;
    FOR(i,8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x,z,k,si
        FOR(i,64) c[i] = (m?m[i]:0)
        u = 1;
    }
}

```

```

FOR(i,20) {
    FOR(j,4) {
        FOR(m,4) t[m] = x[(5*j+4*m)%16];
        t[1] ^= L32(t[0]+t[3], 7);
        t[2] ^= L32(t[1]+t[0], 9);
        t[3] ^= L32(t[2]+t[1], 13);
        t[0] ^= L32(t[3]+t[2], 18);
        FOR(m,4) w[4*j+(j+m)%4] = t[m];
    }
    FOR(m,16) x[m] = w[m];
}

FOR(i,16) st32(out + 4 * i,x[i] + y[i]);
return 0;
}

```

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8 *c,
const u8 *m,u64 b,const u8 *n,const u8 *k)
{
    u8 z[16],x[64];
    u32 u,i;
    if (!b) return 0;
    FOR(i,16) z[i] = 0;
    FOR(i,8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x,z,k,sigma);
        FOR(i,64) c[i] = (m?m[i]:0) ^ x[i];
        u = 1;
    }
}

```

5

```

0) {
,4) {
(m,4) t[m] = x[(5*j+4*m)%16];

] ^= L32(t[0]+t[3], 7);
] ^= L32(t[1]+t[0], 9);
] ^= L32(t[2]+t[1], 13);
] ^= L32(t[3]+t[2], 18);

(m,4) w[4*j+(j+m)%4] = t[m];

,16) x[m] = w[m];

6) st32(out + 4 * i, x[i] + y[i]);
0;

```

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8 *c,
const u8 *m,u64 b,const u8 *n,const u8 *k)
{
    u8 z[16],x[64];
    u32 u,i;
    if (!b) return 0;
    FOR(i,16) z[i] = 0;
    FOR(i,8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x,z,k,sigma);
        FOR(i,64) c[i] = (m?m[i]:0) ^ x[i];
        u = 1;

```

6

```

for (
    u +=
    z[i]
    u >
}
b -=
c +=
if (m)
}
if (b) -
crypto
FOR(i
}
return
}

```

5

```

= x[(5*j+4*m)%16];
...+t[3], 7);
...+t[0], 9);
...+t[1], 13);
...+t[2], 18);
...-(j+m)%4] = t[m];
...w[m];
...+ 4 * i, x[i] + y[i]);

```

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8 *c,
const u8 *m, u64 b, const u8 *n, const u8 *k)
{
    u8 z[16], x[64];
    u32 u, i;
    if (!b) return 0;
    FOR(i, 16) z[i] = 0;
    FOR(i, 8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x, z, k, sigma);
        FOR(i, 64) c[i] = (m?m[i]:0) ^ x[i];
        u = 1;
    }
}

```

6

```

for (i = 8; i < 16)
    u += (u32) z[i];
z[i] = u;
u >>= 8;
}
b -= 64;
c += 64;
if (m) m += 64;
}
if (b) {
    crypto_core_salsa20(x, z, k, sigma);
    FOR(i, b) c[i] = (m?m[i]:0) ^ x[i];
}
return 0;
}

```

5

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8 *c,
const u8 *m,u64 b,const u8 *n,const u8 *k)
{
    u8 z[16],x[64];
    u32 u,i;
    if (!b) return 0;
    FOR(i,16) z[i] = 0;
    FOR(i,8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x,z,k,sigma);
        FOR(i,64) c[i] = (m?m[i]:0) ^ x[i];
        u = 1;

```

6

```

        for (i = 8;i < 16;++i) {
            u += (u32) z[i];
            z[i] = u;
            u >>= 8;
        }
        b -= 64;
        c += 64;
        if (m) m += 64;
    }
    if (b) {
        crypto_core_salsa20(x,z,k,si
        FOR(i,b) c[i] = (m?m[i]:0) ^
    }
    return 0;
}

```

```

static const u8 sigma[16]
= "expand 32-byte k";

int crypto_stream_salsa20_xor(u8 *c,
const u8 *m,u64 b,const u8 *n,const u8 *k)
{
    u8 z[16],x[64];
    u32 u,i;
    if (!b) return 0;
    FOR(i,16) z[i] = 0;
    FOR(i,8) z[i] = n[i];
    while (b >= 64) {
        crypto_core_salsa20(x,z,k,sigma);
        FOR(i,64) c[i] = (m?m[i]:0) ^ x[i];
        u = 1;

```

```

        for (i = 8;i < 16;++i) {
            u += (u32) z[i];
            z[i] = u;
            u >>= 8;
        }
        b -= 64;
        c += 64;
        if (m) m += 64;
    }
    if (b) {
        crypto_core_salsa20(x,z,k,sigma);
        FOR(i,b) c[i] = (m?m[i]:0) ^ x[i];
    }
    return 0;
}

```


6

```

const u8 sigma[16]
    "32-byte k";

crypto_stream_salsa20_xor(u8 *c,
    *m,u64 b,const u8 *n,const u8 *k)
    ,x[64];
;
return 0;
) z[i] = 0;
) z[i] = n[i];
o >= 64) {
o_core_salsa20(x,z,k,sigma);
,64) c[i] = (m?m[i]:0) ^ x[i];
;

```

7

```

for (i = 8;i < 16;++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}

b -= 64;
c += 64;

if (m) m += 64;
}

if (b) {
    crypto_core_salsa20(x,z,k,sigma);
    FOR(i,b) c[i] = (m?m[i]:0) ^ x[i];
}

return 0;
}

```

Next step
For each
compile
and see

6

```

a[16]
...
crypto_core_salsa20_xor(u8 *c,
...
const u8 *n, const u8 *k)
...
];
...
crypto_core_salsa20(x, z, k, sigma);
...
(m?m[i]:0) ^ x[i];

```

```

for (i = 8; i < 16; ++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}

b -= 64;
c += 64;
if (m) m += 64;
}

if (b) {
    crypto_core_salsa20(x, z, k, sigma);
    FOR(i, b) c[i] = (m?m[i]:0) ^ x[i];
}

return 0;
}

```

7

Next step in analysis
 For each target CPU
 compile the simple
 and see how fast it

6

```

*c,
nst u8 *k)

sigma);
^ x[i];
}

```

```

for (i = 8;i < 16;++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}

b -= 64;

c += 64;

if (m) m += 64;
}

if (b) {
    crypto_core_salsa20(x,z,k,sigma);
    FOR(i,b) c[i] = (m?m[i]:0) ^ x[i];
}

return 0;
}

```

7

Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

```
for (i = 8; i < 16; ++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}

b -= 64;

c += 64;

if (m) m += 64;
}

if (b) {
    crypto_core_salsa20(x, z, k, sigma);
    FOR(i, b) c[i] = (m ? m[i] : 0) ^ x[i];
}

return 0;
}
```

Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

```
for (i = 8; i < 16; ++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}
b -= 64;
c += 64;
if (m) m += 64;
}
if (b) {
    crypto_core_salsa20(x, z, k, sigma);
    FOR(i, b) c[i] = (m ? m[i] : 0) ^ x[i];
}
return 0;
}
```

Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

In compiler writer's fantasy world,
the analysis now ends.

```

for (i = 8; i < 16; ++i) {
    u += (u32) z[i];
    z[i] = u;
    u >>= 8;
}
b -= 64;
c += 64;
if (m) m += 64;
}
if (b) {
    crypto_core_salsa20(x, z, k, sigma);
    FOR(i, b) c[i] = (m ? m[i] : 0) ^ x[i];
}
return 0;
}

```

Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

In compiler writer's fantasy world,
the analysis now ends.

“We come so close to optimal on most architectures that we can't do much more without using NP complete algorithms instead of heuristics. We can only try to get little niggles here and there where the heuristics get slightly wrong answers.”

```

i = 8; i < 16; ++i) {
    = (u32) z[i];
    ] = u;
    >= 8;

    64;
    64;
    ) m += 64;

    {
    o_core_salsa20(x, z, k, sigma);
    , b) c[i] = (m ? m[i] : 0) ^ x[i];

    0;

```

7

Next step in analysis:
 For each target CPU,
 compile the simple code,
 and see how fast it is.

In compiler writer's fantasy world,
 the analysis now ends.

“We come so close to optimal on
 most architectures that we can't
 do much more without using NP
 complete algorithms instead of
 heuristics. We can only try to
 get little niggles here and there
 where the heuristics get
 slightly wrong answers.”

8

Reality i




```

s; ++i) {
;

```

Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

In compiler writer's fantasy world,
the analysis now ends.

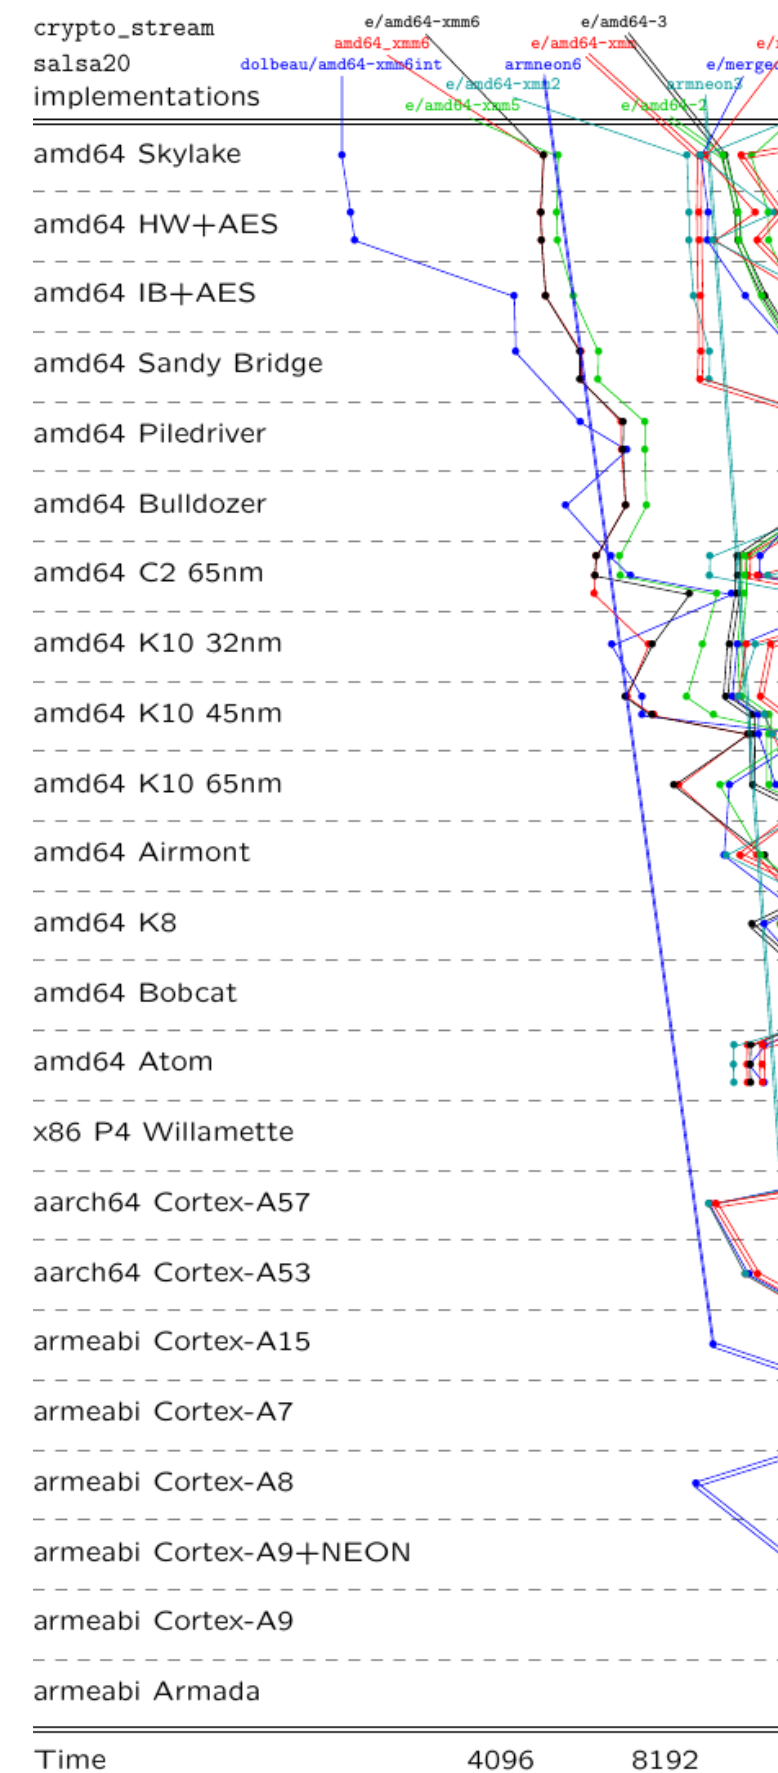
“We come so close to optimal on most architectures that we can't do much more without using NP complete algorithms instead of heuristics. We can only try to get little niggles here and there where the heuristics get slightly wrong answers.”

```

salsa20(x, z, k, sigma);
(m?m[i]:0) ^ x[i];

```

Reality is more complex



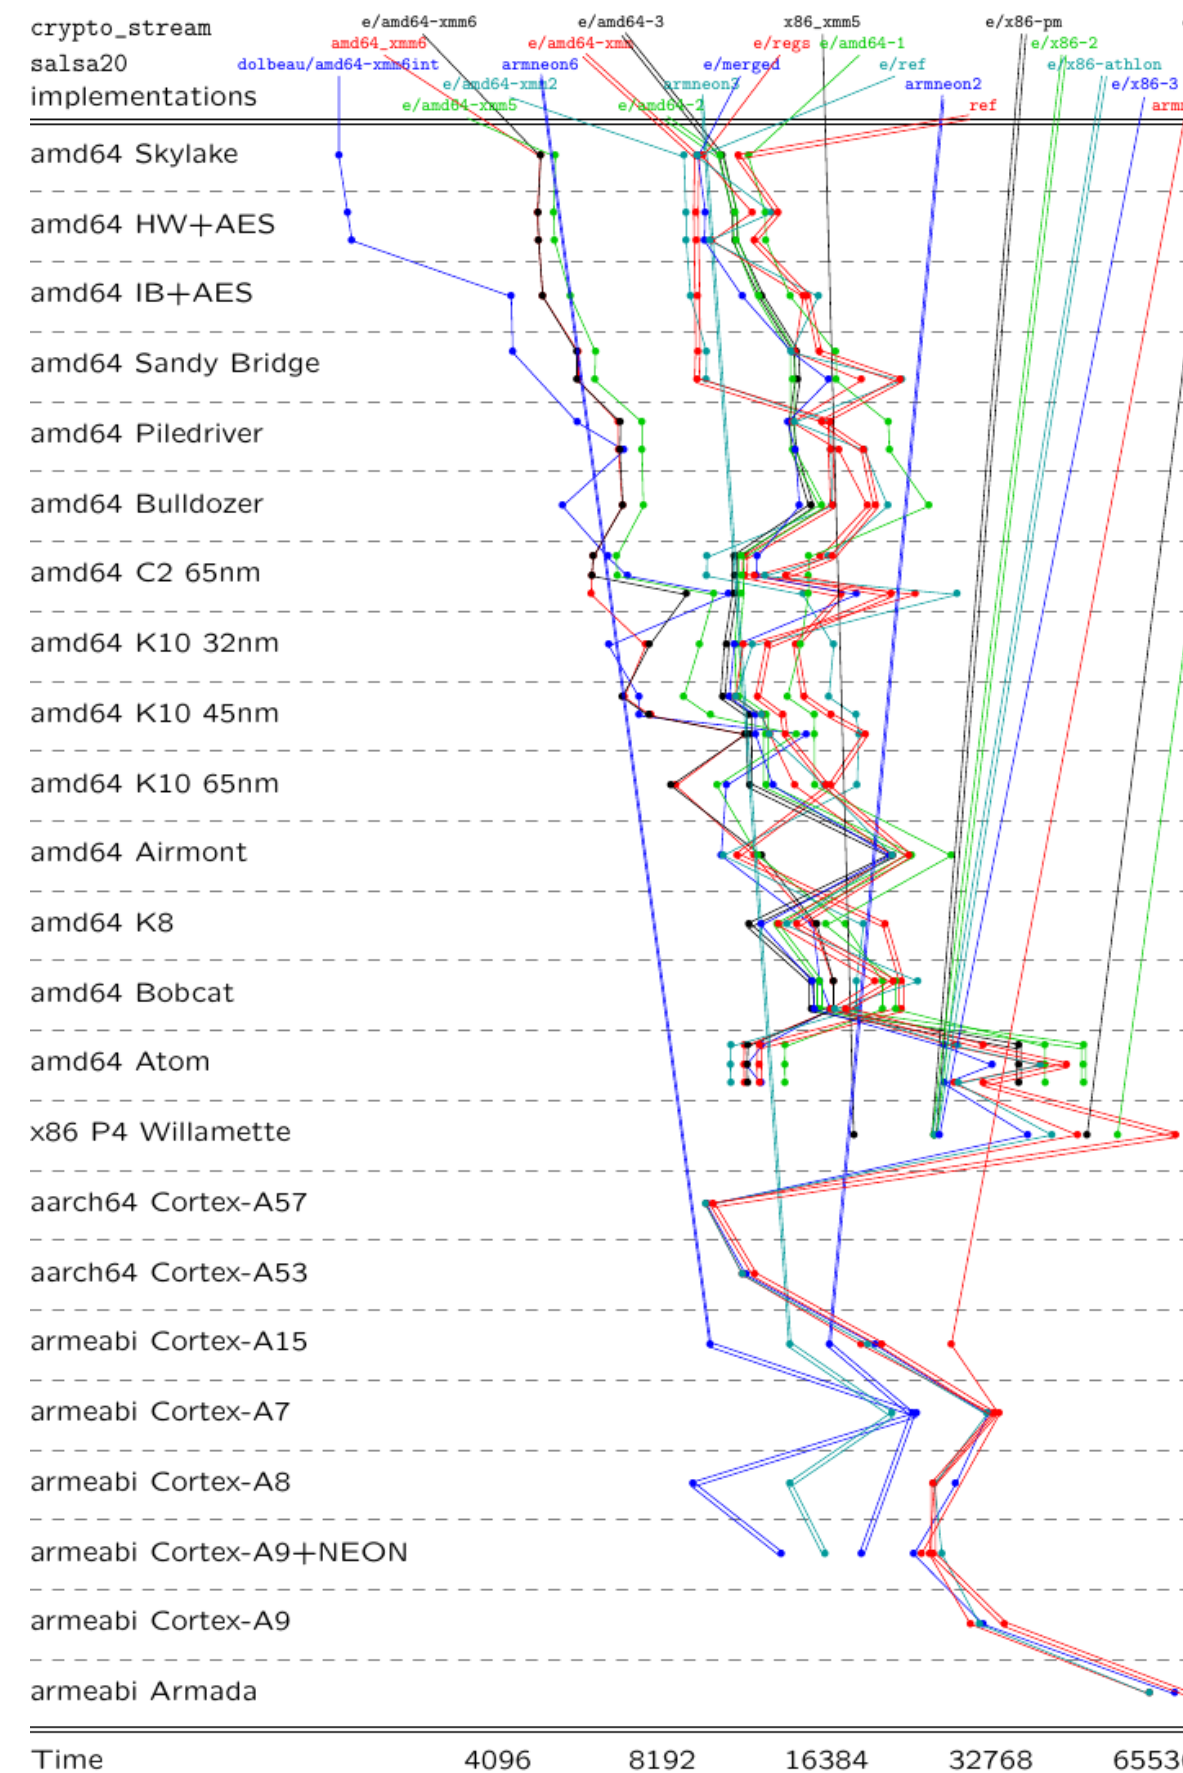
Next step in analysis:
For each target CPU,
compile the simple code,
and see how fast it is.

In compiler writer's fantasy world,
the analysis now ends.

“We come so close to optimal on most architectures that we can't do much more without using NP complete algorithms instead of heuristics. We can only try to get little niggles here and there where the heuristics get slightly wrong answers.”

```
gma);  
x[i];
```

Reality is more complicated:



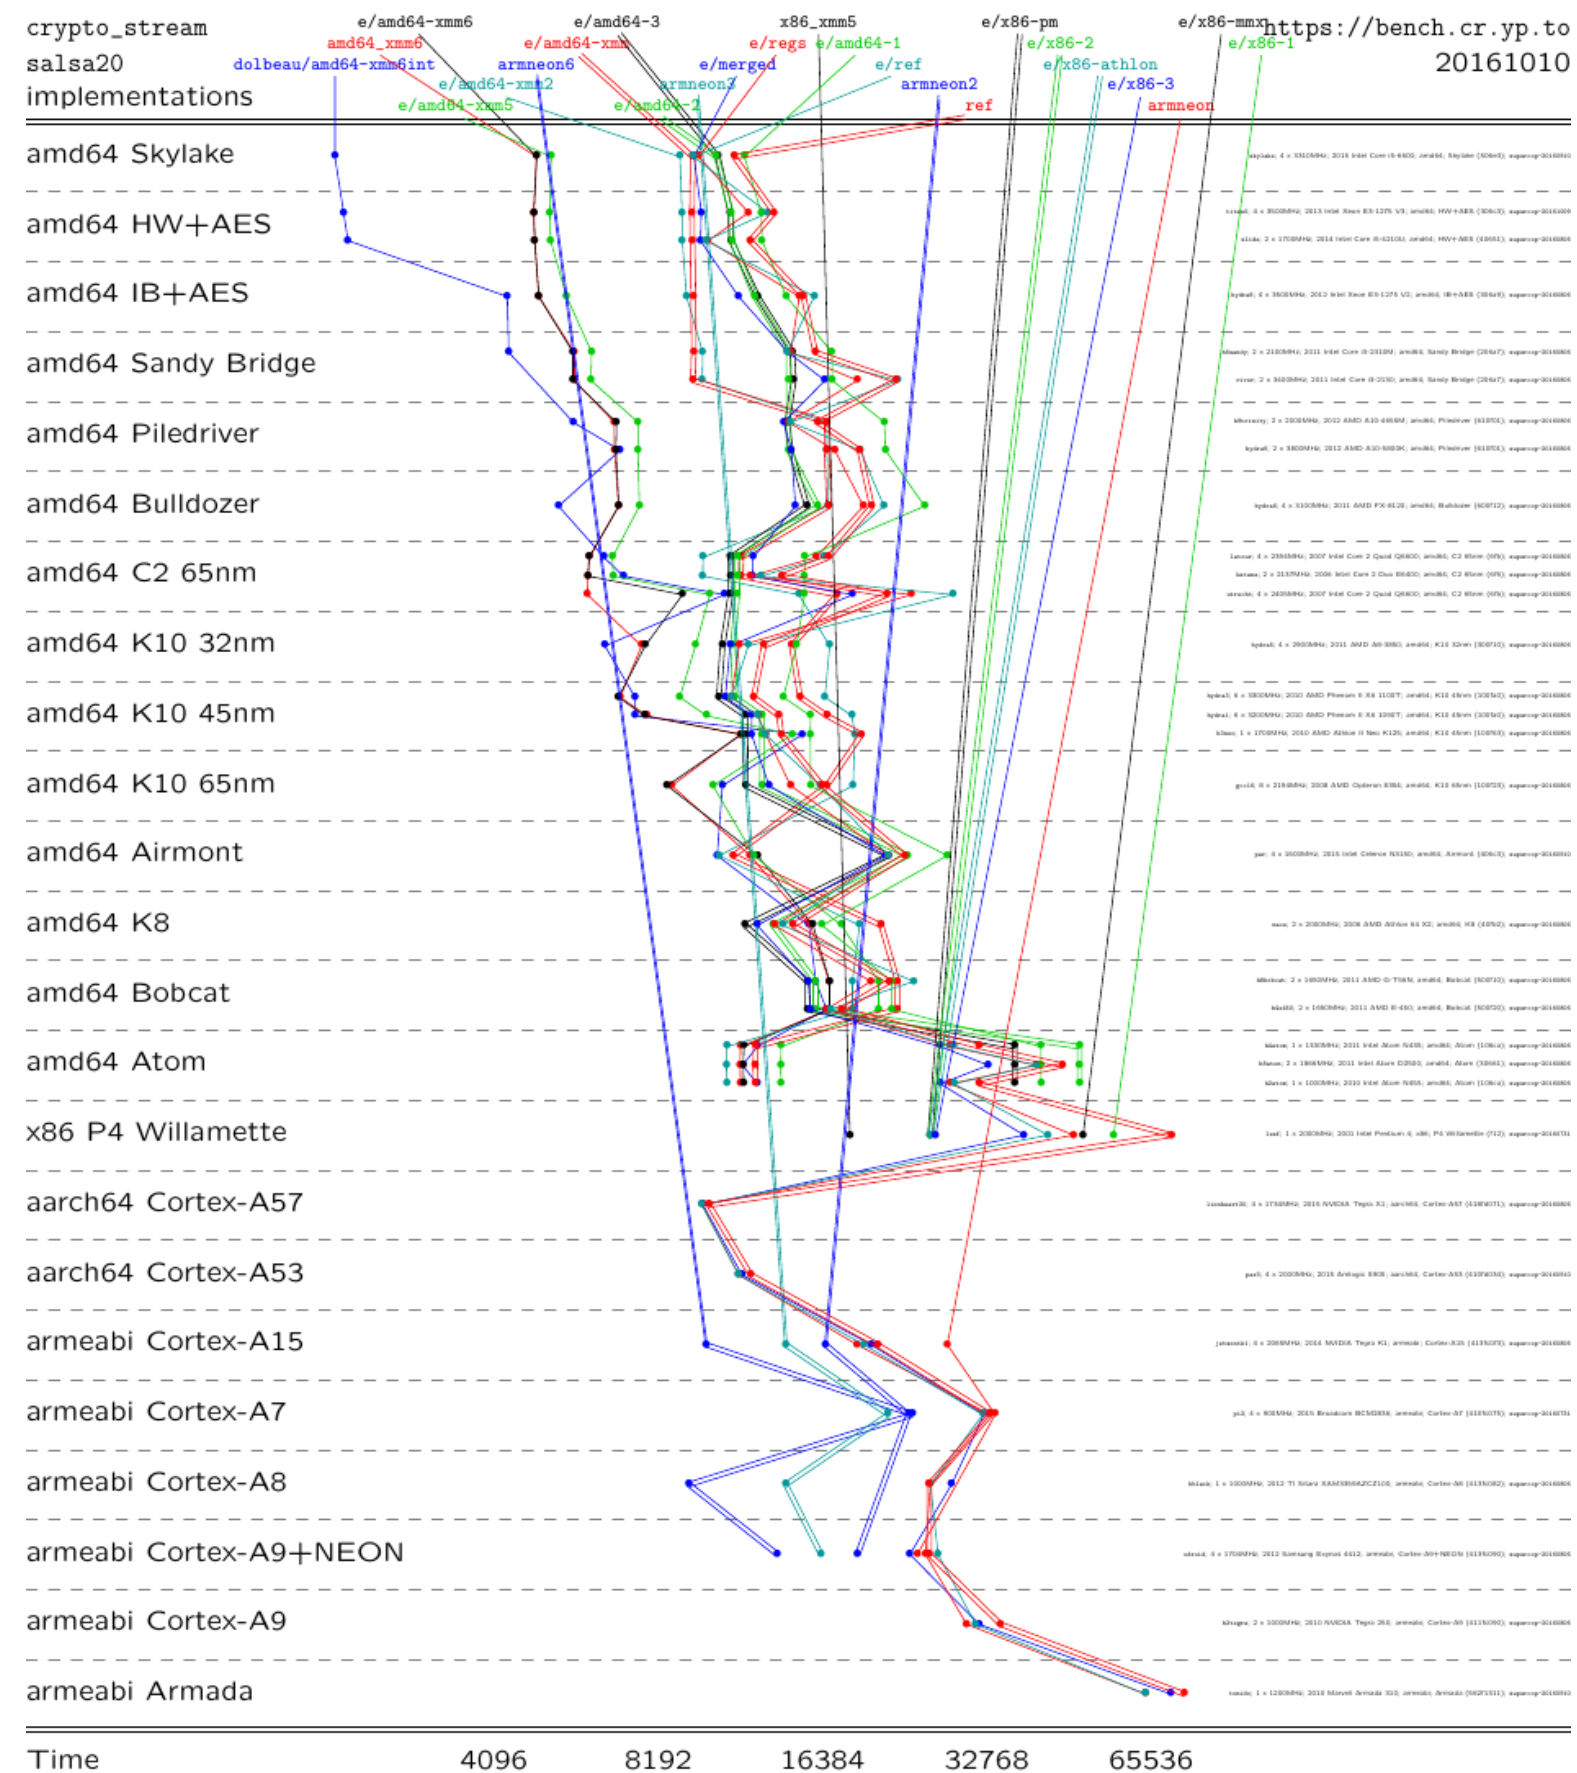
Next step in analysis:

For each target CPU,
compile the simple code,
and see how fast it is.

In compiler writer's fantasy world,
the analysis now ends.

“We come so close to optimal on most architectures that we can't do much more without using NP complete algorithms instead of heuristics. We can only try to get little niggles here and there where the heuristics get slightly wrong answers.”

Reality is more complicated:



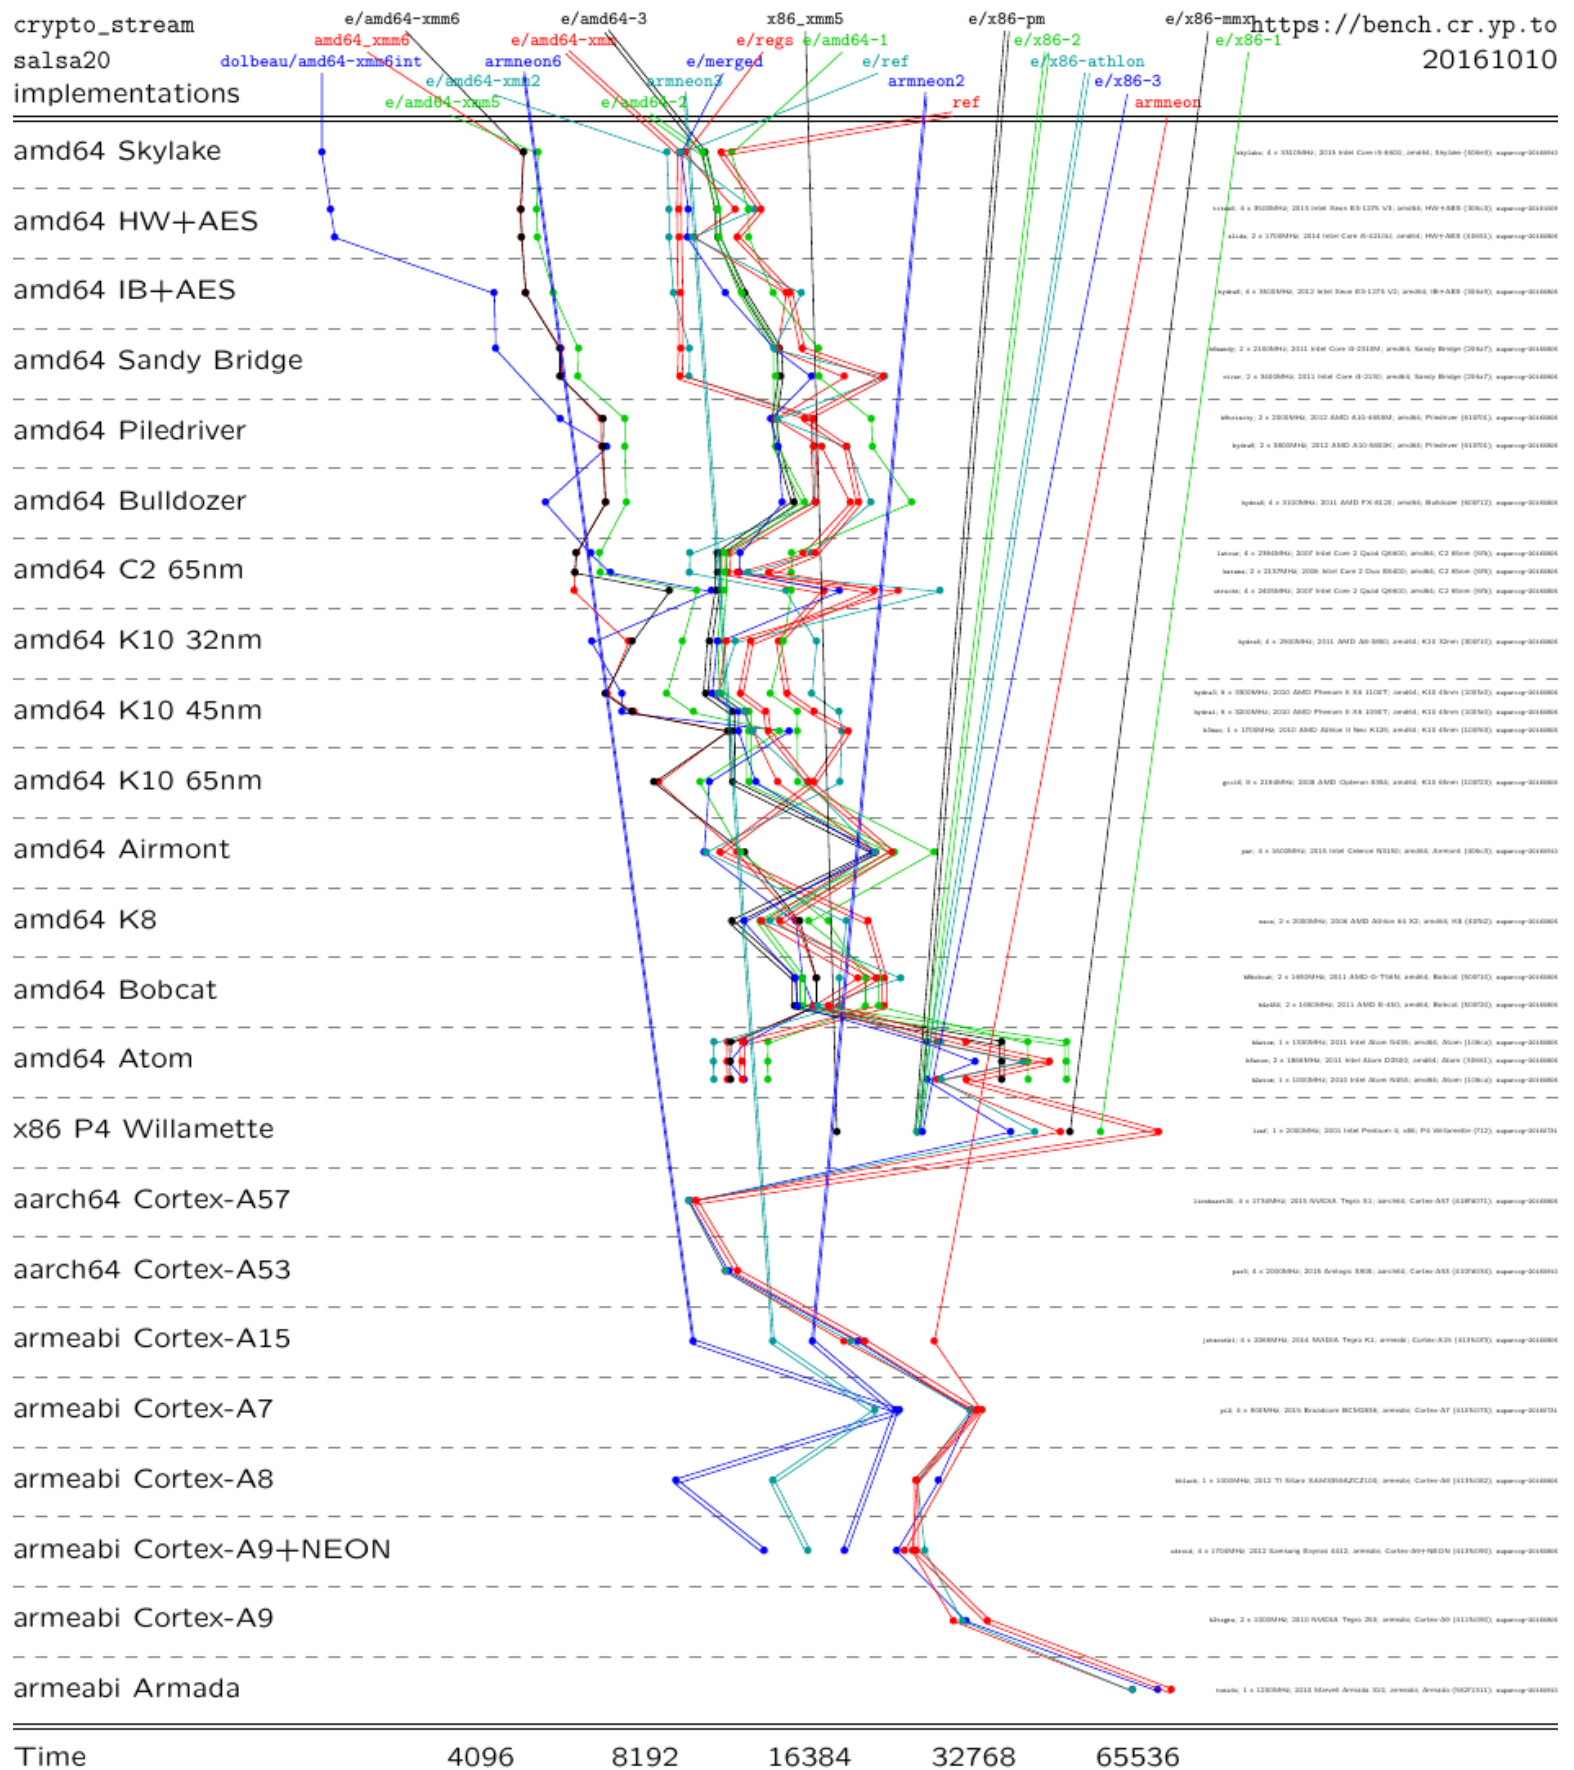
step in analysis:

target CPU,
the simple code,
how fast it is.

compiler writer's fantasy world,
analysis now ends.

come so close to optimal on
architectures that we can't
do more without using NP
complete algorithms instead of
heuristics. We can only try to
fix the niggles here and there
until the heuristics get
wrong answers."

Reality is more complicated:

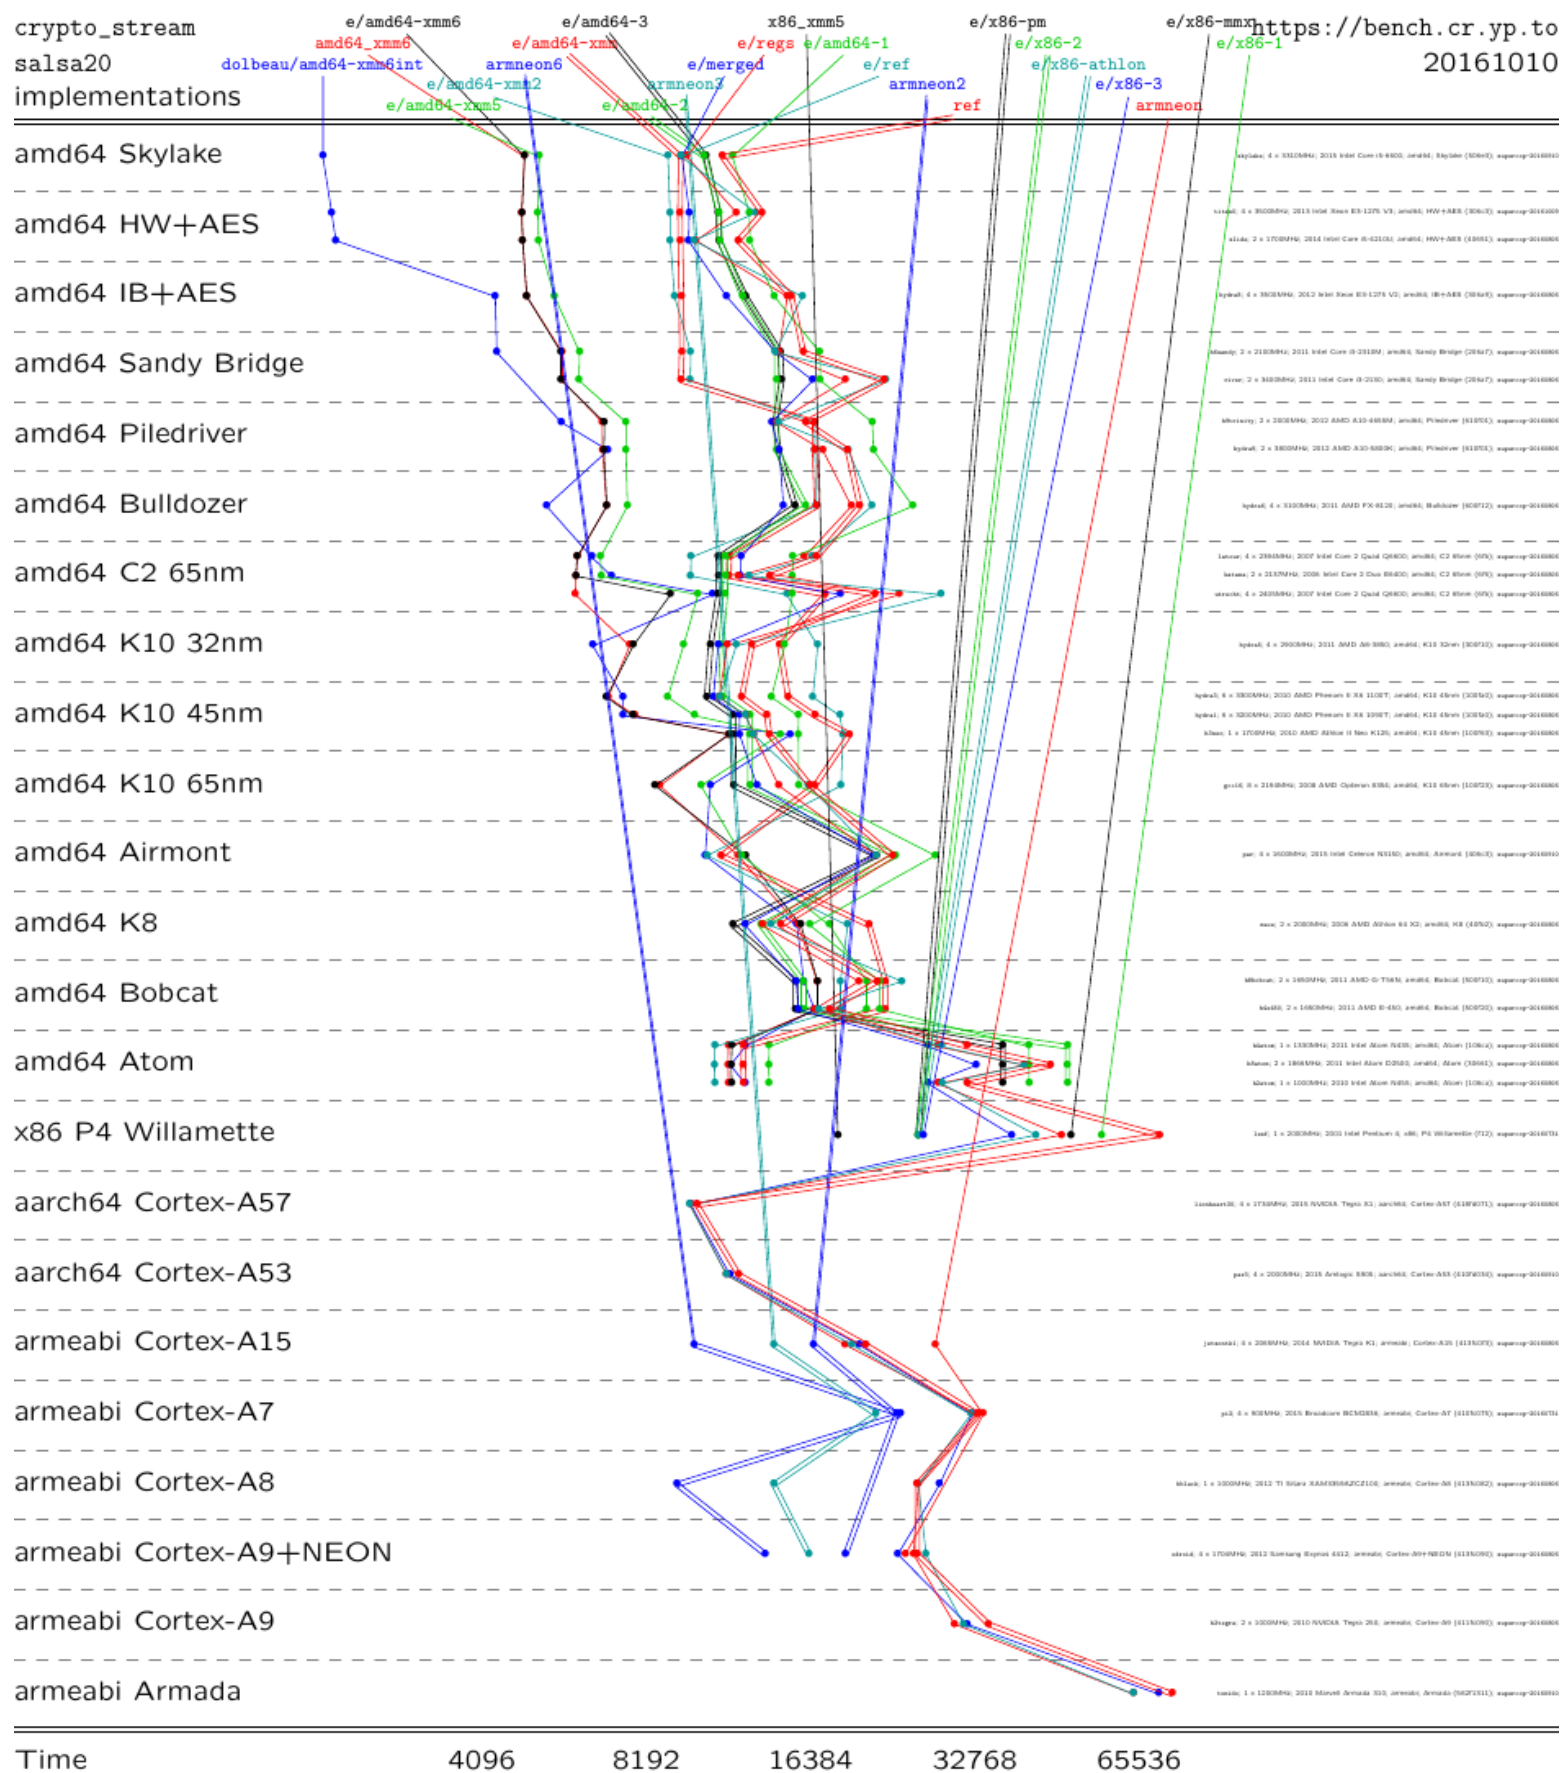


SUPERCOOL
includes
of 563 c
>20 imp
Haswell:
impleme
gcc -O3
is 6.15x
Salsa20
merged
with "m
optimiza
compiler

sis:
 PU,
 e code,
 t is.
 s fantasy world,
 nds.
 e to optimal on
 s that we can't
 hout using NP
 ns instead of
 n only try to
 ere and there
 cs get
 wers."

8

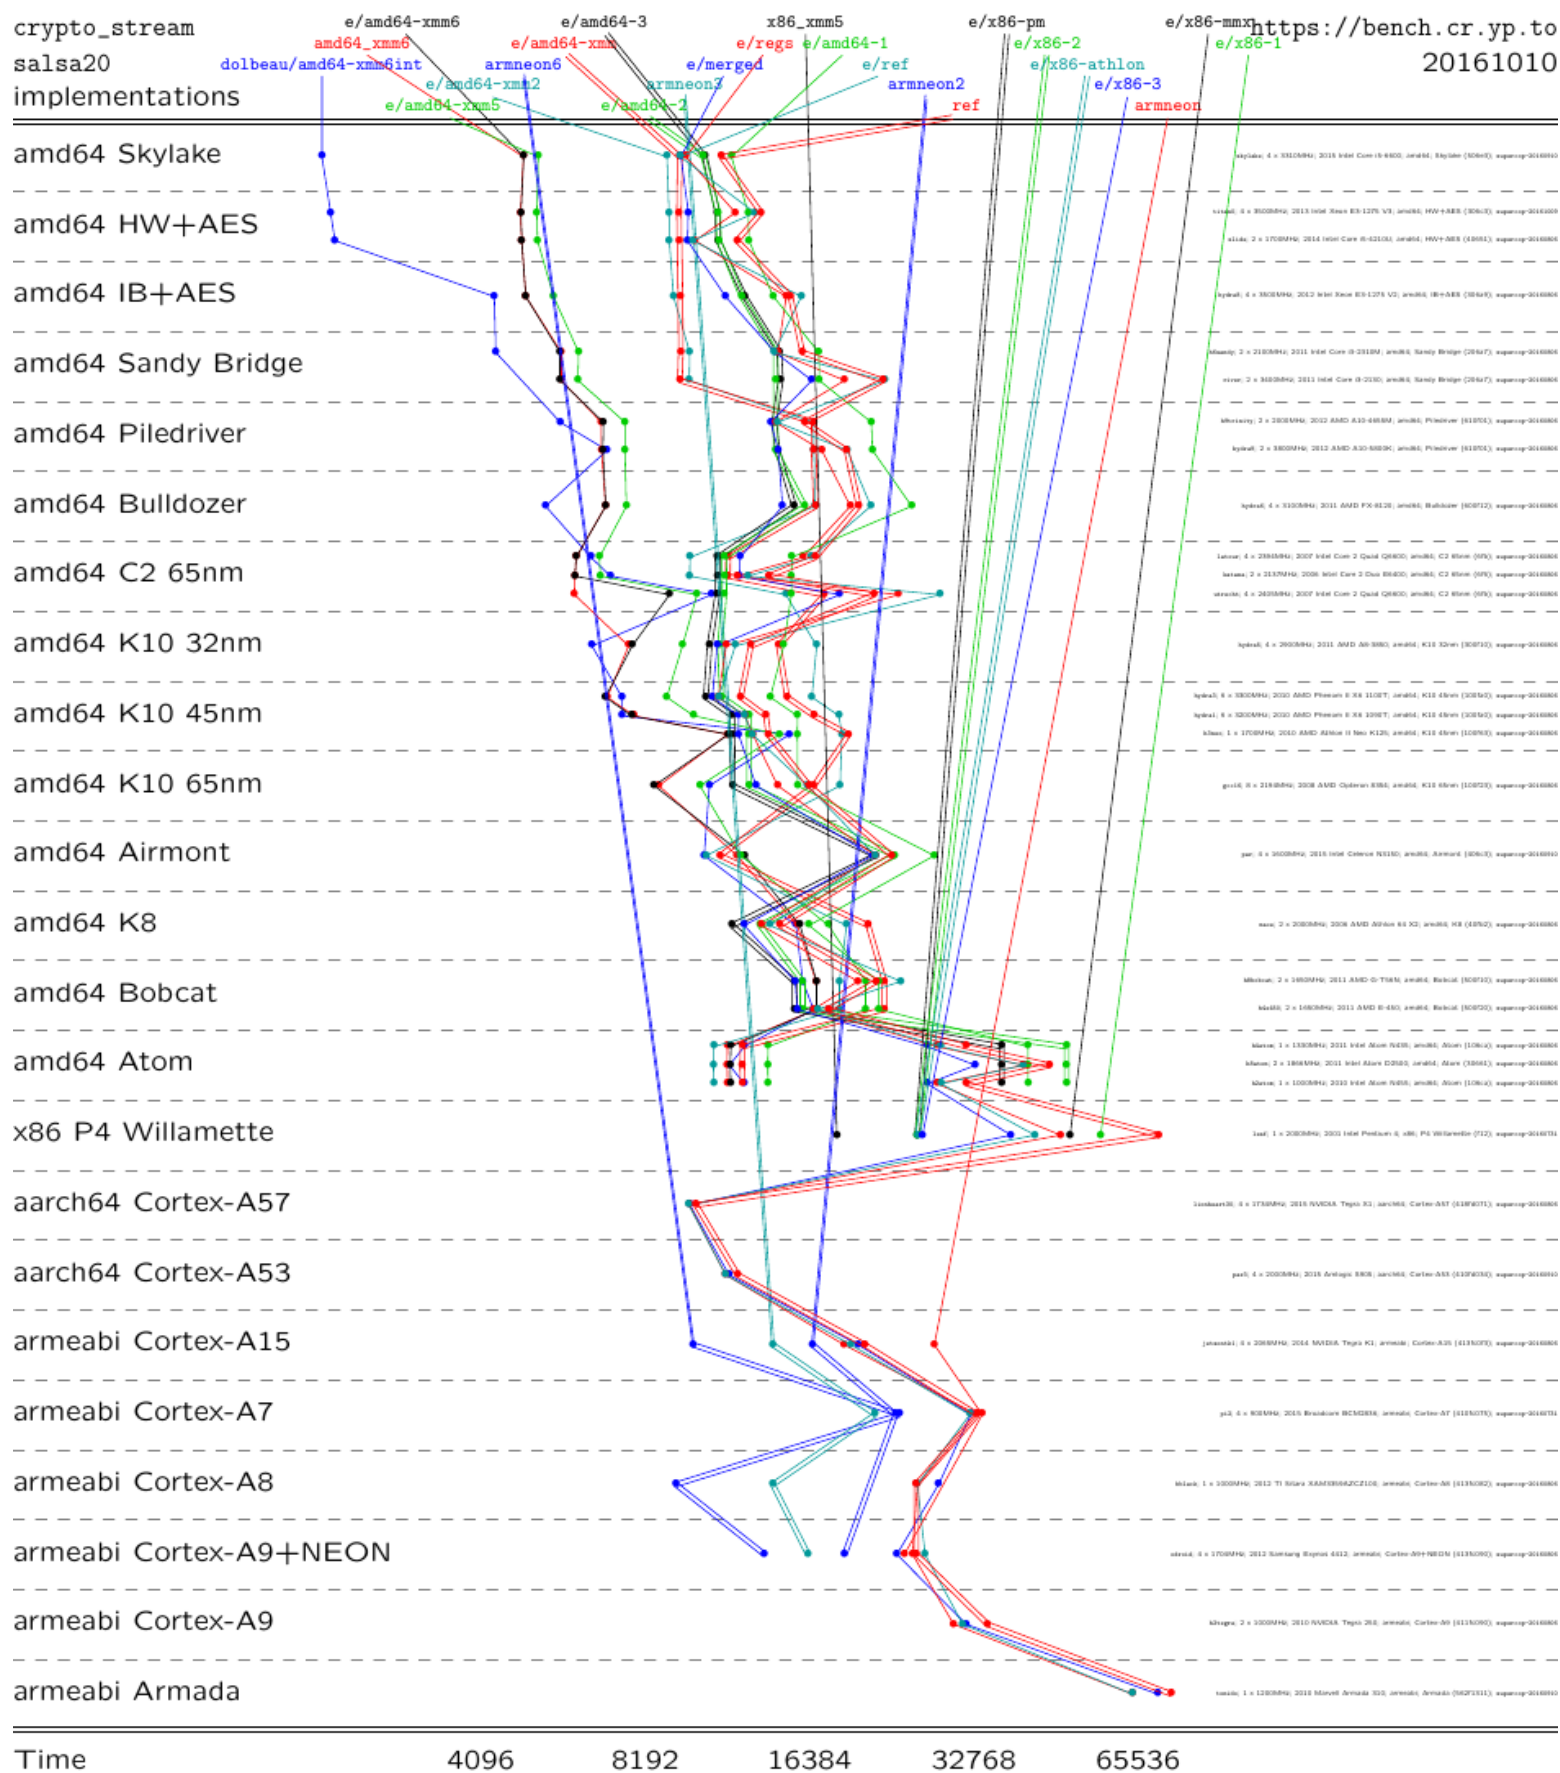
Reality is more complicated:



9

SUPERCOP bench
 includes 2064 imp
 of 563 cryptograph
 >20 implementati
 Haswell: Reasonab
 implementation co
 gcc -O3 -fomit-
 is 6.15x slower th
 Salsa20 implement
 merged implement
 with "machine-ind
 optimizations and
 compiler options:

Reality is more complicated:



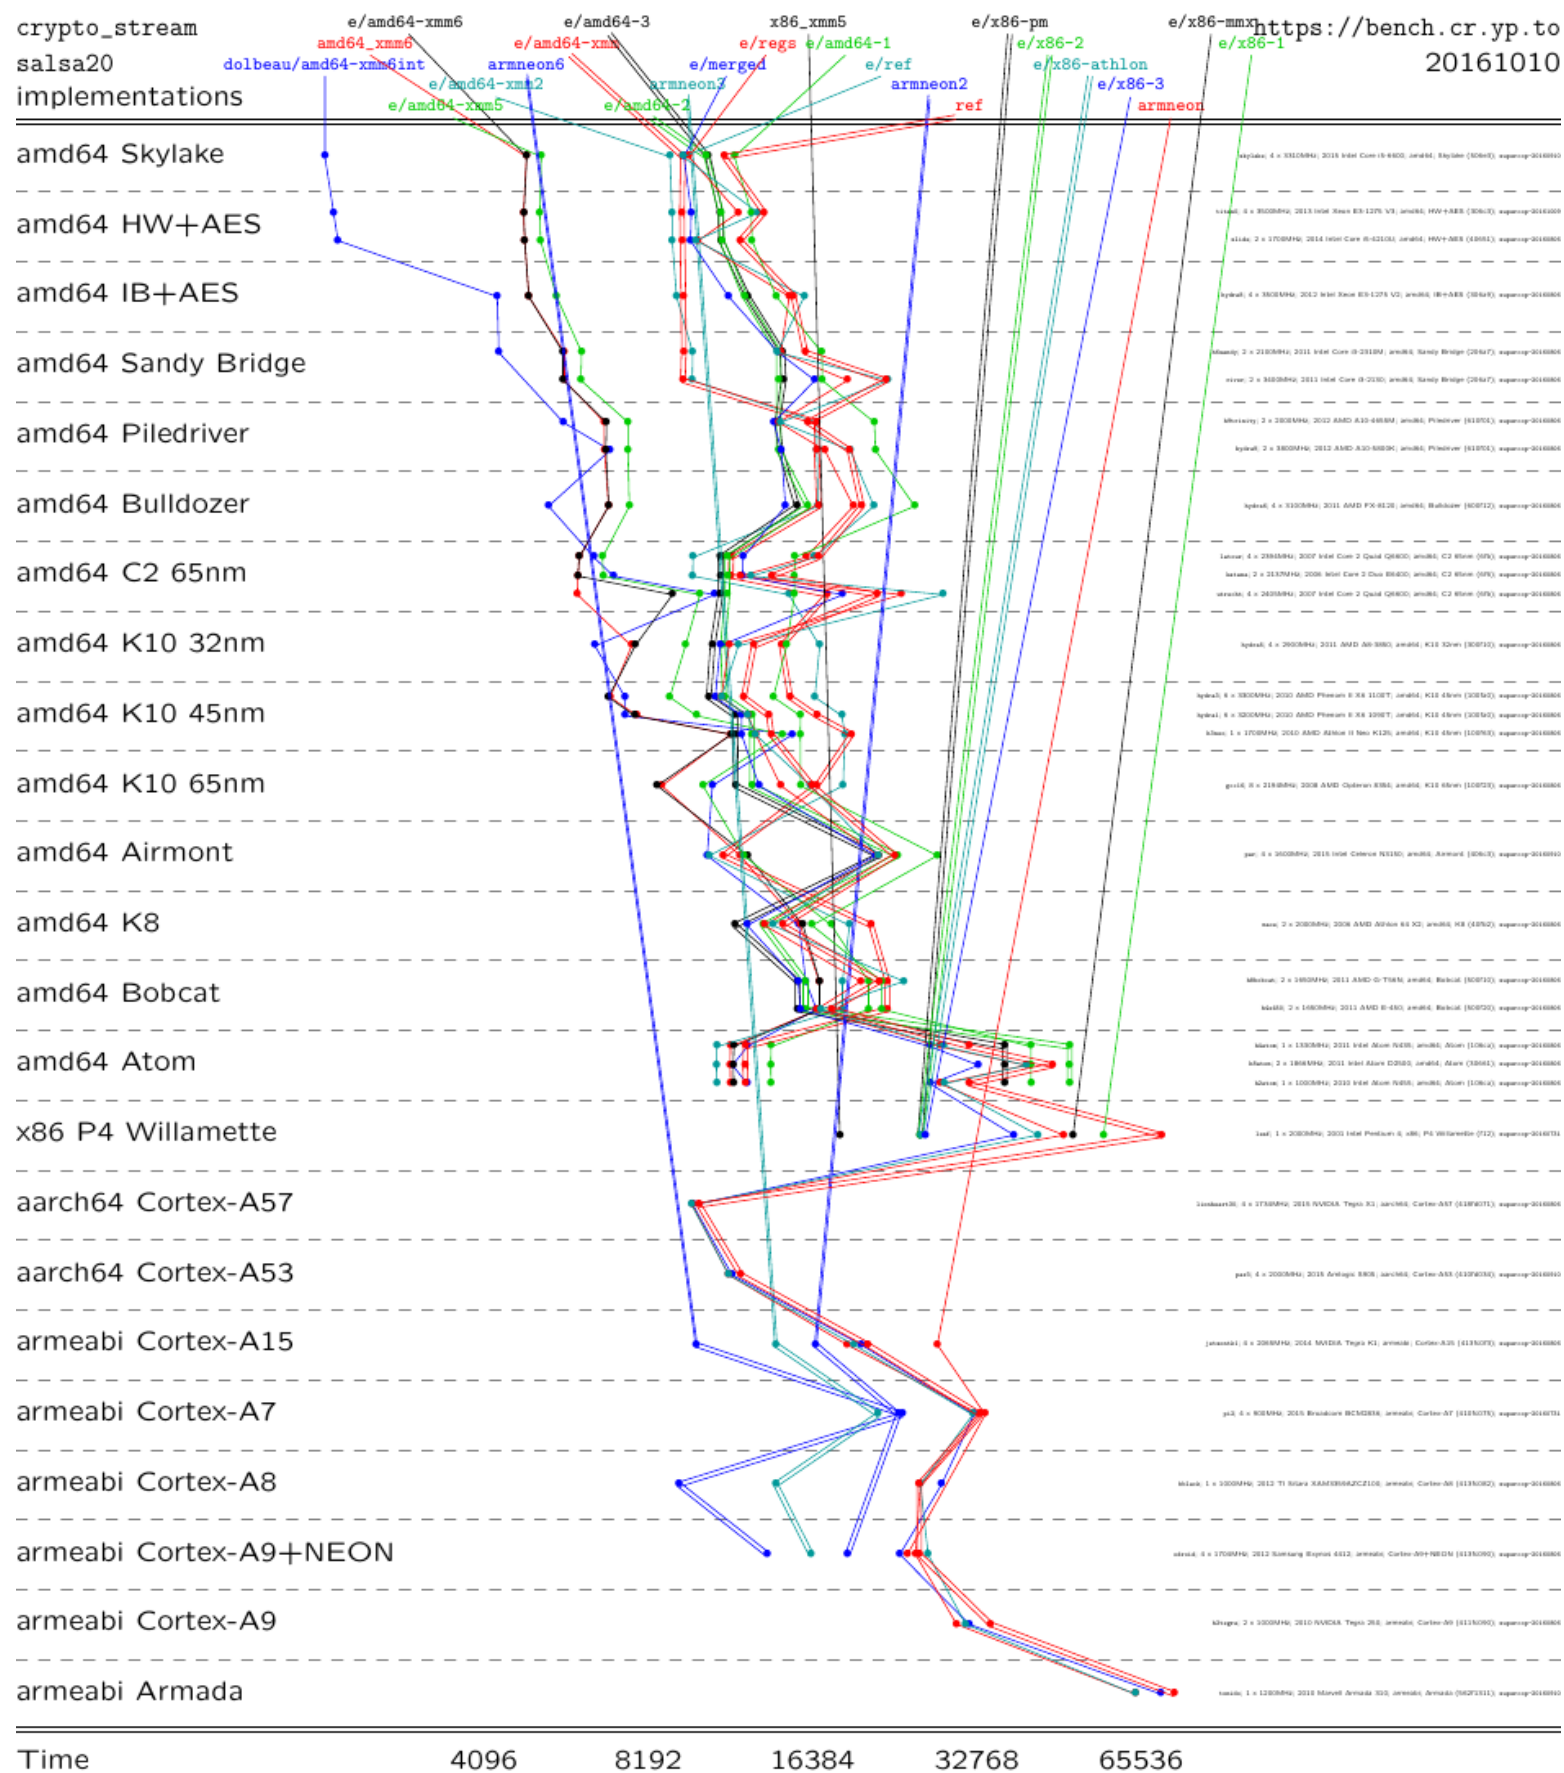
SUPERCOP benchmarking tool includes 2064 implementations of 563 cryptographic primitives >20 implementations of Salsa20

Haswell: Reasonably simple implementation compiled with gcc -O3 -fomit-frame-pointer is 6.15x slower than fastest Salsa20 implementation.

merged implementation with "machine-independent" optimizations and best of 12 compiler options: 4.52x slower

world,
can't
NP
of
to
ere

Reality is more complicated:

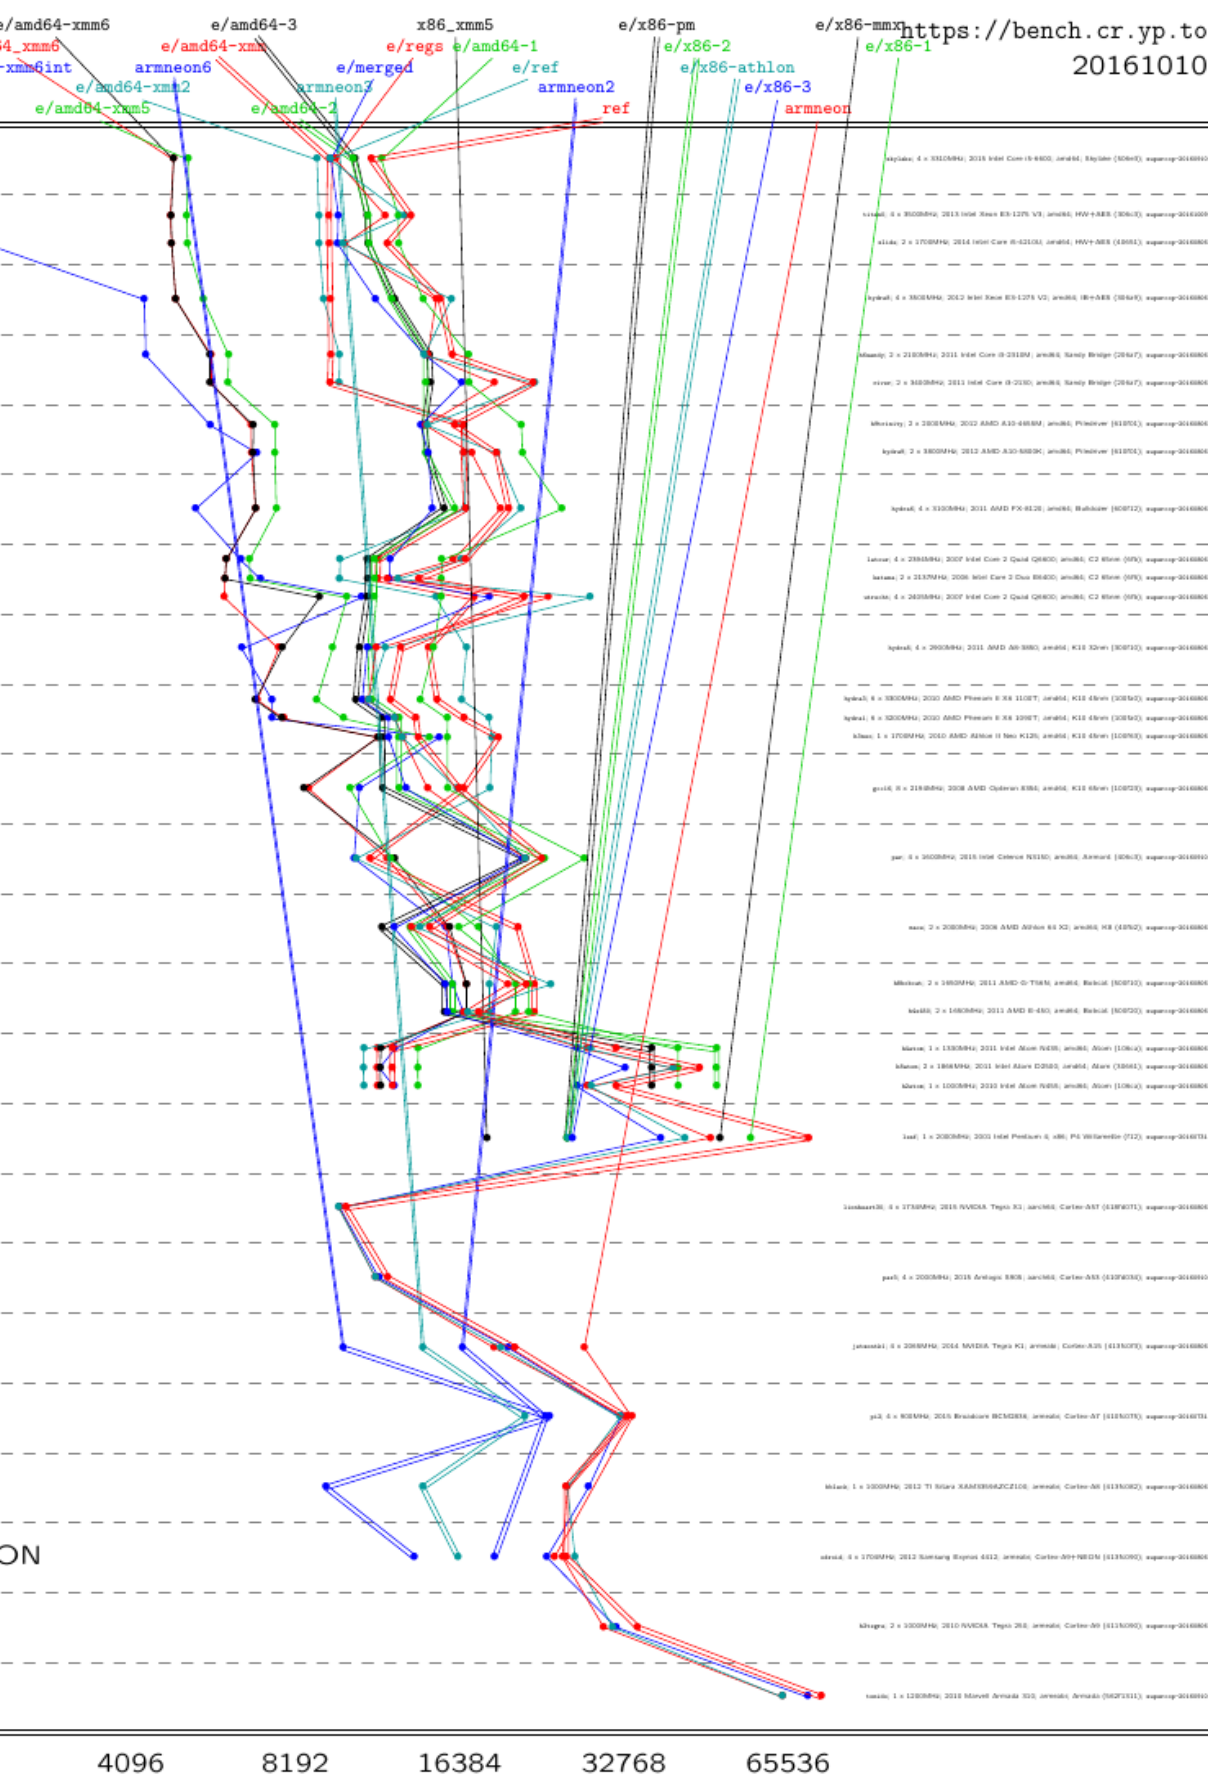


SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

is more complicated:



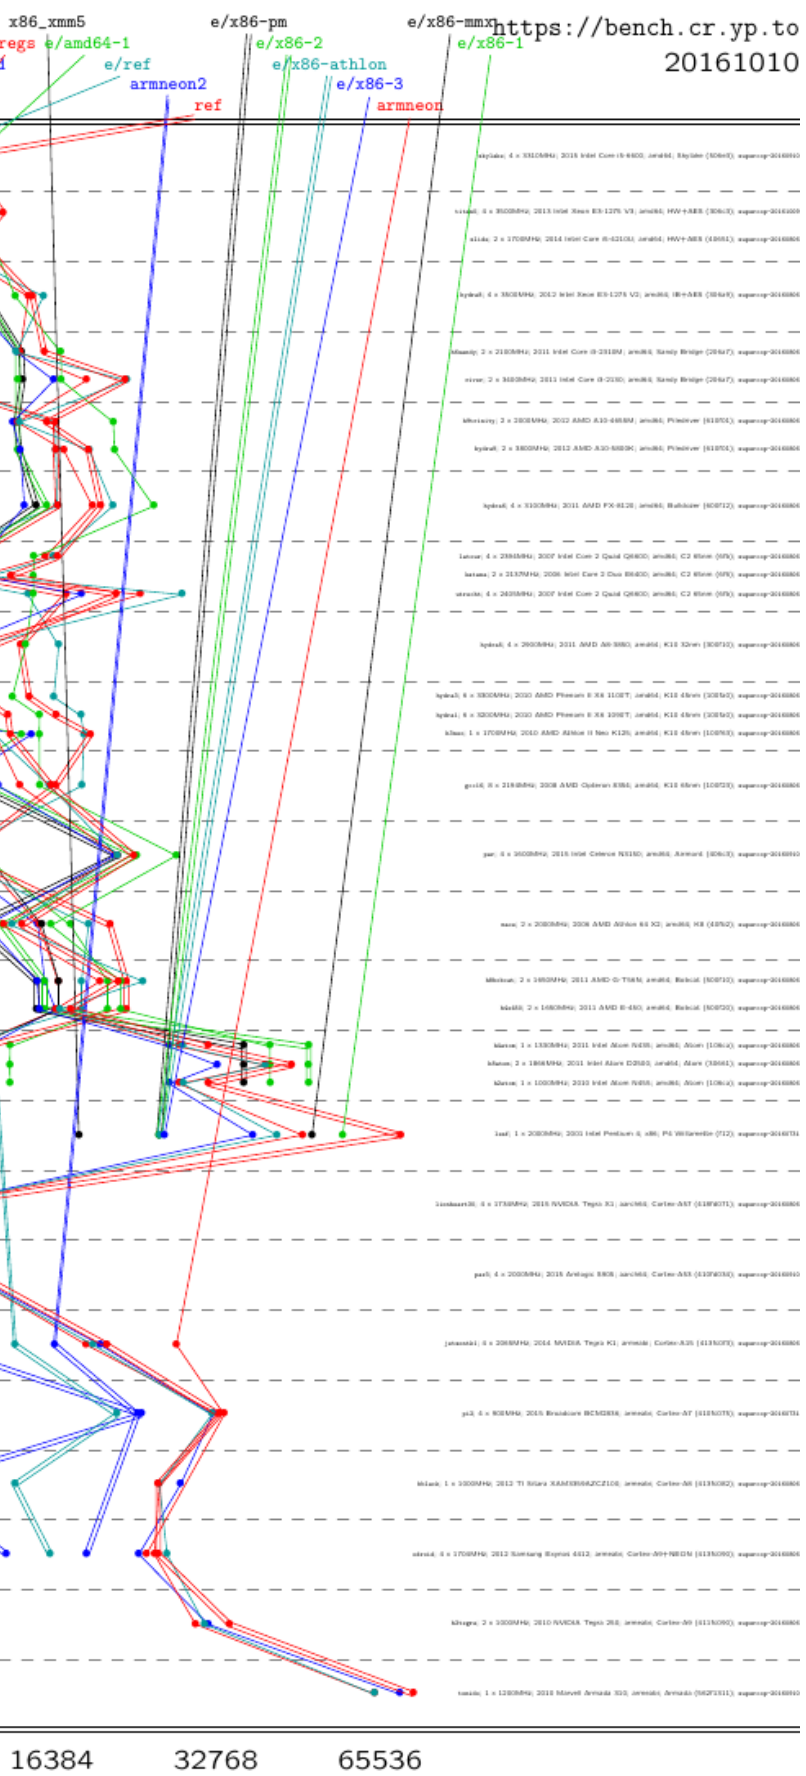
SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with gcc -O3 -fomit-frame-pointer is 6.15x slower than fastest Salsa20 implementation.

merged implementation with "machine-independent" optimizations and best of 121 compiler options: 4.52x slower.

Many m... were dev... to the (c... impleme...

complicated:



SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on to the (currently) implementation fo

neon

4 x 3000MHz	2013 Intel Core i5-4300	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2013 Intel Core i5-4300	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2014 Intel Core i5-4300	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2014 Intel Core i5-4300	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2012 Intel Core i5-3210	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2012 Intel Core i5-3210	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 Intel Core i3-2330	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 Intel Core i3-2330	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2012 AMD A10-7800	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2012 AMD A10-7800	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 AMD FX-8320	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 AMD FX-8320	amd64	OpenSSL (3075)	super-cop/304800
4 x 2000MHz	2007 Intel Core 2 Quad Q9550	amd64	C2 SHA1 (39)	super-cop/304800
2 x 2000MHz	2007 Intel Core 2 Quad Q9550	amd64	C2 SHA1 (39)	super-cop/304800
4 x 2000MHz	2006 Intel Core 2 Duo E6700	amd64	C2 SHA1 (39)	super-cop/304800
2 x 2000MHz	2006 Intel Core 2 Duo E6700	amd64	C2 SHA1 (39)	super-cop/304800
4 x 2000MHz	2007 Intel Core 2 Quad Q9550	amd64	C2 SHA1 (39)	super-cop/304800
2 x 2000MHz	2007 Intel Core 2 Quad Q9550	amd64	C2 SHA1 (39)	super-cop/304800
4 x 2000MHz	2011 AMD A8-3850	amd64	K12 SHA1 (1075)	super-cop/304800
2 x 2000MHz	2011 AMD A8-3850	amd64	K12 SHA1 (1075)	super-cop/304800
4 x 2000MHz	2010 AMD Phenom II X4 960T	amd64	K12 SHA1 (1075)	super-cop/304800
2 x 2000MHz	2010 AMD Phenom II X4 960T	amd64	K12 SHA1 (1075)	super-cop/304800
4 x 1700MHz	2010 AMD Phenom II X4 960T	amd64	K12 SHA1 (1075)	super-cop/304800
2 x 1700MHz	2010 AMD Phenom II X4 960T	amd64	K12 SHA1 (1075)	super-cop/304800
4 x 2100MHz	2008 AMD Catalyst R6	amd64	K12 SHA1 (1075)	super-cop/304800
2 x 2100MHz	2008 AMD Catalyst R6	amd64	K12 SHA1 (1075)	super-cop/304800
4 x 3000MHz	2013 Intel Core i5-3230	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2013 Intel Core i5-3230	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 AMD O-Titan	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 AMD O-Titan	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 AMD R-480	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 AMD R-480	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 Intel Atom D2501	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 Intel Atom D2501	amd64	OpenSSL (3075)	super-cop/304800
4 x 3000MHz	2011 Intel Atom D2501	amd64	OpenSSL (3075)	super-cop/304800
2 x 3000MHz	2011 Intel Atom D2501	amd64	OpenSSL (3075)	super-cop/304800
4 x 2000MHz	2010 Intel Pentium G 690	amd64	OpenSSL (3075)	super-cop/304800
2 x 2000MHz	2010 Intel Pentium G 690	amd64	OpenSSL (3075)	super-cop/304800
4 x 3700MHz	2013 NVIDIA Tegra K1	amd64	Cartes-07 (113407)	super-cop/304800
2 x 3700MHz	2013 NVIDIA Tegra K1	amd64	Cartes-07 (113407)	super-cop/304800
4 x 2000MHz	2013 Analog R88	amd64	Cartes-03 (113403)	super-cop/304800
2 x 2000MHz	2013 Analog R88	amd64	Cartes-03 (113403)	super-cop/304800
4 x 2000MHz	2014 NVIDIA Tegra K1	amd64	Cartes-03 (113403)	super-cop/304800
2 x 2000MHz	2014 NVIDIA Tegra K1	amd64	Cartes-03 (113403)	super-cop/304800
4 x 3000MHz	2013 NVIDIA Tegra K1	amd64	Cartes-07 (113407)	super-cop/304800
2 x 3000MHz	2013 NVIDIA Tegra K1	amd64	Cartes-07 (113407)	super-cop/304800
4 x 3000MHz	2012 TI Sitara XA500	amd64	Cartes-01 (113401)	super-cop/304800
2 x 3000MHz	2012 TI Sitara XA500	amd64	Cartes-01 (113401)	super-cop/304800
4 x 1700MHz	2012 Samsung Exynos 4210	amd64	Cartes-01 (113401)	super-cop/304800
2 x 1700MHz	2012 Samsung Exynos 4210	amd64	Cartes-01 (113401)	super-cop/304800
4 x 3000MHz	2010 NVIDIA Tegra 2	amd64	Cartes-01 (113401)	super-cop/304800
2 x 3000MHz	2010 NVIDIA Tegra 2	amd64	Cartes-01 (113401)	super-cop/304800
4 x 1200MHz	2010 Marvell Armada 300	amd64	Cartes-01 (113401)	super-cop/304800
2 x 1200MHz	2010 Marvell Armada 300	amd64	Cartes-01 (113401)	super-cop/304800

SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives. >20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives.
>20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives.
>20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern.
Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives.
>20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern.
Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

SUPERCOP benchmarking toolkit includes 2064 implementations of 563 cryptographic primitives.
>20 implementations of Salsa20.

Haswell: Reasonably simple ref implementation compiled with `gcc -O3 -fomit-frame-pointer` is $6.15\times$ slower than fastest Salsa20 implementation.

merged implementation with “machine-independent” optimizations and best of 121 compiler options: $4.52\times$ slower.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern.
Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

The software engineer needs fast evaluation of performance.

COP benchmarking toolkit
 2064 implementations
 cryptographic primitives.
 implementations of Salsa20.

Reasonably simple ref
 ntation compiled with
 -fomit-frame-pointer
 slower than fastest
 implementation.

implementation
 machine-independent”
 tions and best of 121
 r options: $4.52\times$ slower.

Many more implementations
 were developed on the way
 to the (currently) fastest
 implementation for this CPU.

This is a common pattern.
 Very fast development cycle:
 modify the implementation,
 check that it still works,
 evaluate its performance.

Results of each evaluation
 guide subsequent modifications.

**The software engineer needs
 fast evaluation of performance.**

The unf
 Slow eva
 is often
 to this o

benchmarking toolkit
 implementations
 of basic primitives.
 versions of Salsa20.

very simple ref
 compiled with
 frame-pointer
 an fastest
 tation.

tation
 dependent”
 best of 121
 4.52× slower.

Many more implementations
 were developed on the way
 to the (currently) fastest
 implementation for this CPU.

This is a common pattern.
 Very fast development cycle:
 modify the implementation,
 check that it still works,
 evaluate its performance.

Results of each evaluation
 guide subsequent modifications.

**The software engineer needs
 fast evaluation of performance.**

The unfortunate re
 Slow evaluation of
 is often a huge ob
 to this optimization

toolkit
ons
ves.
sa20.

ref
th
inter

21
wer.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern. Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

The software engineer needs fast evaluation of performance.

The unfortunate reality:
Slow evaluation of performance is often a huge obstacle to this optimization process.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern.

Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

The software engineer needs fast evaluation of performance.

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern. Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

The software engineer needs fast evaluation of performance.

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

When performance evaluation is too slow, the software engineer has to switch context, and then switching back to optimization produces severe cache misses inside software engineer's brain. ("I'm out of the zone.")

Many more implementations were developed on the way to the (currently) fastest implementation for this CPU.

This is a common pattern. Very fast development cycle: modify the implementation, check that it still works, evaluate its performance.

Results of each evaluation guide subsequent modifications.

The software engineer needs fast evaluation of performance.

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

When performance evaluation is too slow, the software engineer has to switch context, and then switching back to optimization produces severe cache misses inside software engineer's brain. ("I'm out of the zone.")

Often optimization is aborted. ("I'll try some other time.")

more implementations
 developed on the way
 (currently) fastest
 implementation for this CPU.

a common pattern.
 development cycle:
 the implementation,
 that it still works,
 its performance.

of each evaluation
 subsequent modifications.

**Software engineer needs
 evaluation of performance.**

The unfortunate reality:

Slow evaluation of performance
 is often a huge obstacle
 to this optimization process.

When performance evaluation is
 too slow, the software engineer
 has to switch context, and then
 switching back to optimization
 produces severe cache misses
 inside software engineer's brain.
 ("I'm out of the zone.")

Often optimization is aborted.
 ("I'll try some other time.")

Goal of -
 Speed u
 by speed
 "Optimi
 help opt

implementations
 the way
 fastest
 this CPU.

pattern.
 ment cycle:
 mentation,
 works,
 mance.

evaluation
 modifications.

**engineer needs
 f performance.**

The unfortunate reality:

Slow evaluation of performance
 is often a huge obstacle
 to this optimization process.

When performance evaluation is
 too slow, the software engineer
 has to switch context, and then
 switching back to optimization
 produces severe cache misses
 inside software engineer's brain.
 ("I'm out of the zone.")

Often optimization is aborted.
 ("I'll try some other time.")

Goal of this talk:
 Speed up the optimi
 by speeding up be
 "Optimize benchm
 help optimize opti

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

When performance evaluation is too slow, the software engineer has to switch context, and then switching back to optimization produces severe cache misses inside software engineer's brain. ("I'm out of the zone.")

Often optimization is aborted. ("I'll try some other time.")

Goal of this talk:

Speed up the optimization process by speeding up benchmarking

"Optimize benchmarking to help optimize optimization."

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

When performance evaluation is too slow, the software engineer has to switch context, and then switching back to optimization produces severe cache misses inside software engineer's brain. ("I'm out of the zone.")

Often optimization is aborted. ("I'll try some other time.")

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

"Optimize benchmarking to help optimize optimization."

The unfortunate reality:

Slow evaluation of performance is often a huge obstacle to this optimization process.

When performance evaluation is too slow, the software engineer has to switch context, and then switching back to optimization produces severe cache misses inside software engineer's brain. ("I'm out of the zone.")

Often optimization is aborted. ("I'll try some other time.")

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

"Optimize benchmarking to help optimize optimization."

What are the bottlenecks that really need speedups?
Measure the benchmarking process to gain understanding.

"Benchmark benchmarking to help optimize benchmarking."

unfortunate reality:
 evaluation of performance
 a huge obstacle
 optimization process.
 performance evaluation is
 y, the software engineer
 switch context, and then
 g back to optimization
 s severe cache misses
 software engineer's brain.
 t of the zone.")
 optimization is aborted.
 some other time.")

Goal of this talk:
 Speed up the optimization process
 by speeding up benchmarking.
 "Optimize benchmarking to
 help optimize optimization."
 What are the bottlenecks
 that really need speedups?
 Measure the benchmarking
 process to gain understanding.
 "Benchmark benchmarking to
 help optimize benchmarking."

Accessing
 The soft
 on his la
 performa

reality:

F performance
stacale

on process.

e evaluation is

ware engineer

text, and then

optimization

ache misses

gineer's brain.

one.")

n is aborted.

er time.")

Goal of this talk:

Speed up the optimization process
by speeding up benchmarking.

“Optimize benchmarking to
help optimize optimization.”

What are the bottlenecks
that really need speedups?

Measure the benchmarking
process to gain understanding.

“Benchmark benchmarking to
help optimize benchmarking.”

Accessing different

The software engine
on his laptop, but
performance on m

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

“Optimize benchmarking to help optimize optimization.”

What are the bottlenecks that really need speedups?

Measure the benchmarking process to gain understanding.

“Benchmark benchmarking to help optimize benchmarking.”

Accessing different CPUs

The software engineer writes on his laptop, but cares about performance on many more

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

“Optimize benchmarking to help optimize optimization.”

What are the bottlenecks that really need speedups?

Measure the benchmarking process to gain understanding.

“Benchmark benchmarking to help optimize benchmarking.”

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

“Optimize benchmarking to help optimize optimization.”

What are the bottlenecks that really need speedups?
Measure the benchmarking process to gain understanding.

“Benchmark benchmarking to help optimize benchmarking.”

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:
A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Goal of this talk:

Speed up the optimization process by speeding up benchmarking.

“Optimize benchmarking to help optimize optimization.”

What are the bottlenecks that really need speedups?
Measure the benchmarking process to gain understanding.

“Benchmark benchmarking to help optimize benchmarking.”

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:
A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

this talk:

up the optimization process
 including up benchmarking.

ize benchmarking to
 imize optimization.”

e the bottlenecks
 ily need speedups?

the benchmarking
 to gain understanding.

mark benchmarking to
 imize benchmarking.”

Accessing different CPUs

The software engineer writes code
 on his laptop, but cares about
 performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms”
 actually has slower algorithms
 running on faster processors.

Systematic fix: Optimize
 each algorithm, new or old,
 for older and newer processors.

For each

Find a n

copy cod

(assumin

collect n

imization process
enchmarking.

enchmarking to
imization.”

lenecks
peedups?
enchmarking
derstanding.

enchmarking to
enchmarking.”

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

For each target CPU
Find a machine with
copy code to that
(assuming it's on
collect measurements

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

For each target CPU:

Find a machine with that CPU
copy code to that machine
(assuming it's on the Internet)
collect measurements there.

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

For each target CPU:
Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

For each target CPU:

Find a machine with that CPU, copy code to that machine (assuming it's on the Internet), collect measurements there.

But, for security reasons, most machines on the Internet disallow access by default, except access by the owner.

Accessing different CPUs

The software engineer writes code on his laptop, but cares about performance on many more CPUs.

Or at least *should* care!

Surprisingly common failure:

A paper with “faster algorithms” actually has slower algorithms running on faster processors.

Systematic fix: Optimize each algorithm, new or old, for older and newer processors.

For each target CPU:

Find a machine with that CPU, copy code to that machine (assuming it's on the Internet), collect measurements there.

But, for security reasons, most machines on the Internet disallow access by default, except access by the owner.

Solution #1: Each software engineer buys each CPU.

This is expensive at high end, time-consuming at low end.

Using different CPUs

Software engineer writes code on laptop, but cares about performance on many more CPUs.

What *should* care!

Common failure:

Software with “faster algorithms”

but has slower algorithms

on faster processors.

Typical fix: Optimize

algorithm, new or old,

and newer processors.

For each target CPU:

Find a machine with that CPU, copy code to that machine (assuming it's on the Internet), collect measurements there.

But, for security reasons, most machines on the Internet disallow access by default, except access by the owner.

Solution #1: Each software engineer buys each CPU.

This is expensive at high end, time-consuming at low end.

Solution

Poor cov

at CPUs

neer writes code
cares about
any more CPUs.

care!

on failure:
ter algorithms”
r algorithms
processors.

optimize
ew or old,
er processors.

For each target CPU:
Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

But, for security reasons,
most machines on the Internet
disallow access by default,
except access by the owner.

Solution #1: Each software
engineer buys each CPU.

This is expensive at high end,
time-consuming at low end.

Solution #2: Ama
Poor coverage of C

14

s code
ut
CPUs.

ms”
ns

ors.

For each target CPU:
Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

But, for security reasons,
most machines on the Internet
disallow access by default,
except access by the owner.

Solution #1: Each software
engineer buys each CPU.

This is expensive at high end,
time-consuming at low end.

15

Solution #2: Amazon.
Poor coverage of CPUs.

For each target CPU:
Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

But, for security reasons,
most machines on the Internet
disallow access by default,
except access by the owner.

Solution #1: Each software
engineer buys each CPU.

This is expensive at high end,
time-consuming at low end.

Solution #2: Amazon.
Poor coverage of CPUs.

For each target CPU:
Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

But, for security reasons,
most machines on the Internet
disallow access by default,
except access by the owner.

Solution #1: Each software
engineer buys each CPU.

This is expensive at high end,
time-consuming at low end.

Solution #2: Amazon.
Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.
Coverage of CPUs is better
but not good enough for crypto.
Usual goals are OS coverage
and architecture coverage.

For each target CPU:

Find a machine with that CPU,
copy code to that machine
(assuming it's on the Internet),
collect measurements there.

But, for security reasons,
most machines on the Internet
disallow access by default,
except access by the owner.

Solution #1: Each software
engineer buys each CPU.

This is expensive at high end,
time-consuming at low end.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.
Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying "Are you
willing to run this code?"

Slow; unreliable; scales badly.

target CPU:
 machine with that CPU,
 de to that machine
 ng it's on the Internet),
 measurements there.

security reasons,
 machines on the Internet
 access by default,
 ccess by the owner.

#1: Each software
 buys each CPU.
 expensive at high end,
 nsuming at low end.

Solution #2: Amazon.
 Poor coverage of CPUs.

Solution #3: Compile farms,
 such as GCC Compile Farm.
 Coverage of CPUs is better
 but not good enough for crypto.
 Usual goals are OS coverage
 and architecture coverage.

Solution #4: Figure out who
 has the right machines. (How?)
 Send email saying "Are you
 willing to run this code?"
 Slow; unreliable; scales badly.

Solution
 "Can I h
 Saves tim

CPU:
with that CPU,
machine
(the Internet),
there.
reasons,
the Internet
default,
the owner.
software
CPU.
at high end,
low end.

Solution #2: Amazon.
Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.
Coverage of CPUs is better
but not good enough for crypto.
Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)
Send email saying "Are you
willing to run this code?"
Slow; unreliable; scales badly.

Solution #5: Send
"Can I have an ac
Saves time but les

Solution #2: Amazon.
Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.
Coverage of CPUs is better
but not good enough for crypto.
Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)
Send email saying “Are you
willing to run this code?”
Slow; unreliable; scales badly.

Solution #5: Send email say
“Can I have an account?”
Saves time but less reliable.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.

Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying “Are you
willing to run this code?”

Slow; unreliable; scales badly.

Solution #5: Send email saying

“Can I have an account?”

Saves time but less reliable.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.

Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying “Are you
willing to run this code?”

Slow; unreliable; scales badly.

Solution #5: Send email saying

“Can I have an account?”

Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.

Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying “Are you
willing to run this code?”

Slow; unreliable; scales badly.

Solution #5: Send email saying

“Can I have an account?”

Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Good: For each code submission,
one-time centralized audit.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.

Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying “Are you
willing to run this code?”

Slow; unreliable; scales badly.

Solution #5: Send email saying

“Can I have an account?”

Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Good: For each code submission,
one-time centralized audit.

Good: High reliability,
high coverage, built-in tests.

Solution #2: Amazon.

Poor coverage of CPUs.

Solution #3: Compile farms,
such as GCC Compile Farm.

Coverage of CPUs is better
but not good enough for crypto.

Usual goals are OS coverage
and architecture coverage.

Solution #4: Figure out who
has the right machines. (How?)

Send email saying “Are you
willing to run this code?”

Slow; unreliable; scales badly.

Solution #5: Send email saying

“Can I have an account?”

Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Good: For each code submission,
one-time centralized audit.

Good: High reliability,
high coverage, built-in tests.

Bad: Much too slow.

#2: Amazon.

coverage of CPUs.

#3: Compile farms,

GCC Compile Farm.

coverage of CPUs is better

good enough for crypto.

goals are OS coverage

architecture coverage.

#4: Figure out who

owns the right machines. (How?)

Send email saying "Are you

able to run this code?"

Unreliable; scales badly.

Solution #5: Send email saying

"Can I have an account?"

Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized

effort to find machines.

Good: For each code submission,

one-time centralized audit.

Good: High reliability,

high coverage, built-in tests.

Bad: Much too slow.

The eBA

Software

something

Software

sends pa

centraliz

eBACS

integrate

eBACS

new SUP

currently

Amazon.
 CPUs.
 compile farms,
 compile Farm.
 is better
 ough for crypto.
 S coverage
 overage.
 re out who
 hines. (How?)
 "Are you
 code?"
 cales badly.

Solution #5: Send email saying
 "Can I have an account?"
 Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
 effort to find machines.

Good: For each code submission,
 one-time centralized audit.

Good: High reliability,
 high coverage, built-in tests.

Bad: Much too slow.

The eBACS data f
 Software engineer
 something to benc
 Software engineer
 sends package by
 centralized account
 eBACS manager a
 integrates into SU
 eBACS manager b
 new SUPERCOP p
 currently 26-mega

Solution #5: Send email saying
“Can I have an account?”
Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Good: For each code submission,
one-time centralized audit.

Good: High reliability,
high coverage, built-in tests.

Bad: Much too slow.

The eBACS data flow

Software engineer has impl:
something to benchmark.

Software engineer submits impl
sends package by email or (via
centralized account) git push

eBACS manager audits impl
integrates into SUPERCOP.

eBACS manager builds
new SUPERCOP package:
currently 26-megabyte xz.

Solution #5: Send email saying
“Can I have an account?”
Saves time but less reliable.

Solution #6: eBACS.

Good: One-time centralized
effort to find machines.

Good: For each code submission,
one-time centralized audit.

Good: High reliability,
high coverage, built-in tests.

Bad: Much too slow.

The eBACS data flow

Software engineer has impl:
something to benchmark.

Software engineer submits impl:
sends package by email or (with
centralized account) `git push`.

eBACS manager audits impl,
integrates into SUPERCOP.

eBACS manager builds
new SUPERCOP package:
currently 26-megabyte `xz`.

#5: Send email saying
"Have an account?"
... but less reliable.

#6: eBACS.

One-time centralized
... find machines.

For each code submission,
... centralized audit.

High reliability,
... average, built-in tests.

... too slow.

The eBACS data flow

Software engineer has impl:
... something to benchmark.

Software engineer submits impl:
... sends package by email or (with
... centralized account) git push.

eBACS manager audits impl,
... integrates into SUPERCOP.

eBACS manager builds
... new SUPERCOP package:
... currently 26-megabyte xz.

eBACS ...
... and anno

Each ma
... waits un
... is suffici

Each ma
... downloa

SUPERC
... stored o

On a typ
... millions

The eBACS data flow

Software engineer has impl:
something to benchmark.

Software engineer submits impl:
sends package by email or (with
centralized account) `git push`.

eBACS manager audits impl,
integrates into SUPERCOP.

eBACS manager builds
new SUPERCOP package:
currently 26-megabyte `xz`.

eBACS manager u
and announces pa

Each machine ope
waits until the ma
is sufficiently idle.

Each machine ope
downloads SUPER

SUPERCOP scans
stored on disk fro

On a typical high-
millions of files, se

The eBACS data flow

Software engineer has impl:
something to benchmark.

Software engineer submits impl:
sends package by email or (with
centralized account) `git push`.

eBACS manager audits impl,
integrates into SUPERCOP.

eBACS manager builds
new SUPERCOP package:
currently 26-megabyte `xz`.

eBACS manager uploads
and announces package.

Each machine operator
waits until the machine
is sufficiently idle.

Each machine operator
downloads SUPERCOP, runs.

SUPERCOP scans data
stored on disk from previous
On a typical high-end CPU:
millions of files, several GB.

The eBACS data flow

Software engineer has impl:
something to benchmark.

Software engineer submits impl:
sends package by email or (with
centralized account) `git push`.

eBACS manager audits impl,
integrates into SUPERCOP.

eBACS manager builds
new SUPERCOP package:
currently 26-megabyte `xz`.

eBACS manager uploads
and announces package.

Each machine operator
waits until the machine
is sufficiently idle.

Each machine operator
downloads SUPERCOP, runs it.

SUPERCOP scans data
stored on disk from previous runs.
On a typical high-end CPU:
millions of files, several GB.

ACS data flow

the engineer has impl:

ing to benchmark.

the engineer submits impl:

ackage by email or (with

ed account) `git push`.

anager audits impl,

es into SUPERCOP.

anager builds

PERCOP package:

y 26-megabyte `xz`.

18

eBACS manager uploads
and announces package.

Each machine operator
waits until the machine
is sufficiently idle.

Each machine operator
downloads SUPERCOP, runs it.

SUPERCOP scans data
stored on disk from previous runs.

On a typical high-end CPU:
millions of files, several GB.

19

For each
SUPERCOP

SUPERCOP
working
saves res

Typically

SUPERCOP
from thi

700-meg

Machine
`data.gz`

flow

has impl:

chmark.

submits impl:

email or (with

t) git push.

udits impl,

PERCOP.

uilds

package:

byte xz.

eBACS manager uploads
and announces package.

Each machine operator
waits until the machine
is sufficiently idle.

Each machine operator
downloads SUPERCOP, runs it.

SUPERCOP scans data
stored on disk from previous runs.

On a typical high-end CPU:
millions of files, several GB.

For each new impl
SUPERCOP comp

SUPERCOP meas
working compiled
saves results on di

Typically at least a

SUPERCOP collec
from this machine

700-megabyte dat

Machine operator

data.gz, announc

eBACS manager uploads
and announces package.

Each machine operator
waits until the machine
is sufficiently idle.

Each machine operator
downloads SUPERCOP, runs it.

SUPERCOP scans data
stored on disk from previous runs.

On a typical high-end CPU:
millions of files, several GB.

For each new impl-compiler
SUPERCOP compiles+tests

SUPERCOP measures each
working compiled impl,
saves results on disk.

Typically at least an hour.

SUPERCOP collects all data
from this machine, typically
700-megabyte data.gz.

Machine operator uploads
data.gz, announces it.

eBACS manager uploads and announces package.

Each machine operator waits until the machine is sufficiently idle.

Each machine operator downloads SUPERCOP, runs it.

SUPERCOP scans data stored on disk from previous runs.

On a typical high-end CPU: millions of files, several GB.

For each new impl-compiler pair, SUPERCOP compiles+tests impl.

SUPERCOP measures each working compiled impl, saves results on disk.

Typically at least an hour.

SUPERCOP collects all data from this machine, typically 700-megabyte data.gz.

Machine operator uploads data.gz, announces it.

manager uploads
 ounces package.

machine operator
 til the machine
 ently idle.

machine operator
 ds SUPERCOP, runs it.

COP scans data
 n disk from previous runs.

ypical high-end CPU:
 of files, several GB.

For each new impl-compiler pair,
 SUPERCOP compiles+tests impl.

SUPERCOP measures each
 working compiled impl,
 saves results on disk.

Typically at least an hour.

SUPERCOP collects all data
 from this machine, typically
 700-megabyte data.gz.

Machine operator uploads
 data.gz, announces it.

eBACS
 data.gz

Databas
 53% cur
 47% arc

For each
 (or for c
 scripts p
 Typically

Web pag
 Under a

uploads
package.

operator
machine

operator
SUPER-COP, runs it.

s data
from previous runs.

end CPU:
several GB.

For each new impl-compiler pair,
SUPER-COP compiles+tests impl.

SUPER-COP measures each
working compiled impl,
saves results on disk.

Typically at least an hour.

SUPER-COP collects all data
from this machine, typically
700-megabyte data.gz.

Machine operator uploads
data.gz, announces it.

eBACS manager c
data.gz into cent

Database currently
53% current uncor
47% archives of su

For each new data
(or for cross-cutting
scripts process all
Typically an hour

Web pages are reg
Under an hour.

For each new impl-compiler pair,
SUPERCOP compiles+tests impl.

SUPERCOP measures each
working compiled impl,
saves results on disk.

Typically at least an hour.

SUPERCOP collects all data
from this machine, typically
700-megabyte data.gz.

Machine operator uploads
data.gz, announces it.

eBACS manager copies
data.gz into central database

Database currently uses 500
53% current uncompressed
47% archives of superseded

For each new data.gz
(or for cross-cutting updates)
scripts process all results.

Typically an hour per machine

Web pages are regenerated.
Under an hour.

For each new impl-compiler pair,
SUPERCOP compiles+tests impl.

SUPERCOP measures each
working compiled impl,
saves results on disk.

Typically at least an hour.

SUPERCOP collects all data
from this machine, typically
700-megabyte data.gz.

Machine operator uploads
data.gz, announces it.

eBACS manager copies
data.gz into central database.

Database currently uses 500GB:
53% current uncompressed data,
47% archives of superseded data.

For each new data.gz
(or for cross-cutting updates):
scripts process all results.
Typically an hour per machine.

Web pages are regenerated.
Under an hour.

a new impl-compiler pair,
COP compiles+tests impl.

COP measures each
compiled impl,
results on disk.

y at least an hour.

COP collects all data
s machine, typically
gabyte data.gz.

e operator uploads
z, announces it.

eBACS manager copies
data.gz into central database.

Database currently uses 500GB:
53% current uncompressed data,
47% archives of superseded data.

For each new data.gz
(or for cross-cutting updates):
scripts process all results.
Typically an hour per machine.

Web pages are regenerated.
Under an hour.

In progress

New data

All impls

Some m

measure

“publish

does for

All comp

All check

All meas

All table

-compiler pair,
files+tests impl.

ures each
impl,
sk.

an hour.

cts all data
, typically
a.gz.

uploads
ces it.

eBACS manager copies
data.gz into central database.

Database currently uses 500GB:
53% current uncompressed data,
47% archives of superseded data.

For each new data.gz
(or for cross-cutting updates):
scripts process all results.
Typically an hour per machine.

Web pages are regenerated.
Under an hour.

In progress: SUPE

New database stor

All impls ever subr

Some metadata no

measurements. Bu

“publish results” f

does force new me

All compiled impls

All checksums of c

All measurements.

All tables, graphs,

pair,
impl.

eBACS manager copies
data.gz into central database.

Database currently uses 500GB:
53% current uncompressed data,
47% archives of superseded data.

For each new data.gz
(or for cross-cutting updates):
scripts process all results.

Typically an hour per machine.

Web pages are regenerated.

Under an hour.

In progress: SUPERCOP 2

New database stored central

All impls ever submitted.

Some metadata not affecting
measurements. But turning
“publish results” for an impl
does force new measurement

All compiled impls.

All checksums of outputs.

All measurements.

All tables, graphs, etc.

eBACS manager copies

data.gz into central database.

Database currently uses 500GB:

53% current uncompressed data,

47% archives of superseded data.

For each new data.gz

(or for cross-cutting updates):

scripts process all results.

Typically an hour per machine.

Web pages are regenerated.

Under an hour.

In progress: SUPERCOP 2

New database stored centrally:

All impls ever submitted.

Some metadata not affecting measurements. But turning on “publish results” for an impl *does* force new measurements.

All compiled impls.

All checksums of outputs.

All measurements.

All tables, graphs, etc.

manager copies
 z into central database.
 e currently uses 500GB:
 rent uncompressed data,
 hives of superseded data.
 a new data.gz
 cross-cutting updates):
 process all results.
 y an hour per machine.
 ges are regenerated.
 n hour.

In progress: SUPERCOP 2

New database stored centrally:

All impls ever submitted.

Some metadata not affecting
 measurements. But turning on
 “publish results” for an impl
does force new measurements.

All compiled impls.

All checksums of outputs.

All measurements.

All tables, graphs, etc.

When ne
 Impl is p
 Each con
 to check
 Each wo
 pushed t
 (when th
 Each me
 immedia
 If impl s
 Measure
 after con

copies
 central database.
 y uses 500GB:
 mpressed data,
 uperseded data.
 a.gz
 ng updates):
 results.
 per machine.
 generated.

In progress: SUPERCOP 2

New database stored centrally:
 All impls ever submitted.
 Some metadata not affecting
 measurements. But turning on
 “publish results” for an impl
does force new measurements.
 All compiled impls.
 All checksums of outputs.
 All measurements.
 All tables, graphs, etc.

When new impl is
 Impl is pushed to
 Each compiled impl
 to checksum mach
 Each working com
 pushed to benchm
 (when they are su
 Each measurement
 immediately to su
 If impl says “publi
 Measurements are
 after comparisons

In progress: SUPERCOP 2

New database stored centrally:

All impls ever submitted.

Some metadata not affecting measurements. But turning on “publish results” for an impl *does* force new measurements.

All compiled impls.

All checksums of outputs.

All measurements.

All tables, graphs, etc.

When new impl is submitted

Impl is pushed to compile servers

Each compiled impl is pushed to checksum machines.

Each working compiled impl pushed to benchmark machines (when they are sufficiently idle)

Each measurement is available immediately to submitter.

If impl says “publish results” Measurements are put online after comparisons are done.

In progress: SUPERCOP 2

New database stored centrally:

All impls ever submitted.

Some metadata not affecting measurements. But turning on “publish results” for an impl *does* force new measurements.

All compiled impls.

All checksums of outputs.

All measurements.

All tables, graphs, etc.

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results” :
Measurements are put online after comparisons are done.

Process: SUPERCOP 2

Database stored centrally:

impls ever submitted.

Metadata not affecting

measurements. But turning on

“publish results” for an impl

triggers new measurements.

Compiled impls.

Checksums of outputs.

Measurements.

Tables, graphs, etc.

22

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed
to checksum machines.

Each working compiled impl is
pushed to benchmark machines
(when they are sufficiently idle).

Each measurement is available
immediately to submitter.

If impl says “publish results”:

Measurements are put online
after comparisons are done.

23

Wait, what?

No more
there's no

Critical i

Can a ro
take over

Or corrup
from oth

ERCOP 2

red centrally:

mitted.

ot affecting

ut turning on

or an impl

measurements.

s.

outputs.

etc.

22

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results”:

Measurements are put online after comparisons are done.

23

Wait, what about

No more central a
there's no time for

Critical integrity c

Can a rogue code
take over the mac

Or corrupt benchm
from other submit

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results”:

Measurements are put online after comparisons are done.

Wait, what about security?

No more central auditing: there's no time for it.

Critical integrity concerns:

Can a rogue code submitter take over the machine?

Or corrupt benchmarks from other submitters?

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results”:

Measurements are put online after comparisons are done.

Wait, what about security?

No more central auditing:
there's no time for it.

Critical integrity concerns:

Can a rogue code submitter
take over the machine?

Or corrupt benchmarks
from other submitters?

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results” :
Measurements are put online after comparisons are done.

Wait, what about security?

No more central auditing:
there's no time for it.

Critical integrity concerns:

Can a rogue code submitter
take over the machine?

Or corrupt benchmarks
from other submitters?

Concerns start before code is
tested and measured: compilers
have bugs, sometimes serious.

When new impl is submitted:

Impl is pushed to compile servers.

Each compiled impl is pushed to checksum machines.

Each working compiled impl is pushed to benchmark machines (when they are sufficiently idle).

Each measurement is available immediately to submitter.

If impl says “publish results”:

Measurements are put online after comparisons are done.

Wait, what about security?

No more central auditing:
there's no time for it.

Critical integrity concerns:

Can a rogue code submitter
take over the machine?

Or corrupt benchmarks
from other submitters?

Concerns start before code is
tested and measured: compilers
have bugs, sometimes serious.

Smaller availability concerns:
e.g., Bitcoin mining.

new impl is submitted:

pushed to compile servers.

compiled impl is pushed

to xsum machines.

working compiled impl is

used to benchmark machines

(when they are sufficiently idle).

Measurement is available

locally to submitter.

Submitter says "publish results":

Results are put online

and comparisons are done.

Wait, what about security?

No more central auditing:

there's no time for it.

Critical integrity concerns:

Can a rogue code submitter

take over the machine?

Or corrupt benchmarks

from other submitters?

Concerns start before code is

tested and measured: compilers

have bugs, sometimes serious.

Smaller availability concerns:

e.g., Bitcoin mining.

SUPERCOMPUTER

OS-level

impl can

cannot f

SUPERCOMPUTER

pool of

each cor

machine

Enforces

for files

in comp

More dif

integrity

tables co

submitted:

compile servers.

impl is pushed

achines.

compiled impl is

mark machines

efficiently idle).

t is available

mitter.

sh results”:

put online

are done.

Wait, what about security?

No more central auditing:

there's no time for it.

Critical integrity concerns:

Can a rogue code submitter

take over the machine?

Or corrupt benchmarks

from other submitters?

Concerns start before code is

tested and measured: compilers

have bugs, sometimes serious.

Smaller availability concerns:

e.g., Bitcoin mining.

SUPERCOP 1 sets

OS-level resource

impl cannot open

cannot fork any pr

SUPERCOP 2 ma

pool of uids and c

each compile serve

machine, benchma

Enforces reasonab

for files legitimately

in compiling an im

More difficult to e

integrity policy for

tables comparing i

Wait, what about security?

No more central auditing:
there's no time for it.

Critical integrity concerns:

Can a rogue code submitter
take over the machine?

Or corrupt benchmarks
from other submitters?

Concerns start before code is
tested and measured: compilers
have bugs, sometimes serious.

Smaller availability concerns:
e.g., Bitcoin mining.

SUPERCOP 1 sets some
OS-level resource limits:
impl cannot open any files,
cannot fork any processes.

SUPERCOP 2 manages
pool of uids and chroot jails
each compile server, checksu
machine, benchmark machine

Enforces reasonable policy
for files legitimately used
in compiling an impl.

More difficult to enforce:
integrity policy for, e.g.,
tables comparing impls.

Wait, what about security?

No more central auditing:
there's no time for it.

Critical integrity concerns:

Can a rogue code submitter
take over the machine?

Or corrupt benchmarks
from other submitters?

Concerns start before code is
tested and measured: compilers
have bugs, sometimes serious.

Smaller availability concerns:
e.g., Bitcoin mining.

SUPERCOP 1 sets some
OS-level resource limits:
impl cannot open any files,
cannot fork any processes.

SUPERCOP 2 manages
pool of uids and chroot jails on
each compile server, checksum
machine, benchmark machine.

Enforces reasonable policy
for files legitimately used
in compiling an impl.

More difficult to enforce:
integrity policy for, e.g.,
tables comparing impls.